A Continuous-Time Model of an Autonomous Aerial Vehicle to Inform and Validate Formal Verification Methods

Murray L. Ireland[a,*], Ruth Hoffmann[a], Alice Miller[a], Gethin Norman[a], Sandor M. Veres[b]

[a] *School of Computing Science, University of Glasgow, Glasgow, UK*
[b] *Autonomous Systems and Robotics Research Group, University of Sheffield, Sheffield, UK*

arXiv:1609.00177v1 [cs.RO] 1 Sep 2016

Abstract

If autonomous vehicles are to be widely accepted, we need to ensure their safe operation. For this reason, verification and validation (V&V) approaches must be developed that are suitable for this domain. Model checking is a formal technique which allows us to exhaustively explore the paths of an abstract model of a system. Using a probabilistic model checker such as PRISM, we may determine properties such as the expected time for a mission, or the probability that a specific mission failure occurs. However, model checking of complex systems is difficult due to the loss of information during abstraction. This is especially so when considering systems such as autonomous vehicles which are subject to external influences. An alternative solution is the use of Monte Carlo simulation to explore the results of a continuous-time model of the system. The main disadvantage of this approach is that the approach is not exhaustive as not all executions of the system are analysed. We are therefore interested in developing a framework for formal verification of autonomous vehicles, using Monte Carlo simulation to inform and validate our symbolic models during the initial stages of development. In this paper, we present a continuous-time model of a quadrotor unmanned aircraft undertaking an autonomous mission. We employ this model in Monte Carlo simulation to obtain specific mission properties which will inform the symbolic models employed in formal verification.

Keywords: control, autonomy, unmanned vehicles, Monte Carlo simulation, quantitative formal verification

1. Introduction

The last few years have seen a huge expansion in the applications and capability of unmanned vehicles. In the military sector, where unmanned vehicles have been in use for decades, efforts are now focussed on increasing the autonomy of such vehicles. That is, increasing their ability to perform tasks with minimal human interference. The commercial sector and hobbyist community are further behind in their development of semi- and fully-autonomous systems. However, autonomous capabilities are just as useful in commercial and civilian applications.

Regardless of the application and their level of autonomy, unmanned vehicles suffer from a number of obstacles which prevent wide scale use. Unmanned aircraft in particular face certification issues related to the dangers posed by a lack of human decision-making in the loop [1]. Vehicles operating in any uncontrolled environment face a number of unpredictable events and random disturbances, whether internal or sourced from the environment. The ability to appropriately respond to such events and manage such disturbances is a crucial part of any vehicle control system. In the push towards autonomy, this control system must remove the human element from the loop entirely.

Verification is the process of ensuring that a system such as an autonomous vehicle performs as intended. That is, given a specific mission, the vehicle is able to realise the mission goals while responding to any number of events, whether internally or externally derived. Many verification methods exist, ranging from the simple and abstract to the computationally-intensive. Systems such as software packages may be modelled and verified using abstract approaches such as formal verification. More complex systems such as autonomous vehicles are less trivial to abstract. This is

*Corresponding author
Email address: Murray.Ireland@glasgow.ac.uk (Murray L. Ireland)

due to the constant variation of the vehicle's behaviour, its control system and any environmental factors with time. The presence of random elements in the system adds further complexity.

Due to the intractability of abstracting a complex system, simulation is often used for verification. While still a form of abstraction, continuous- or discrete-time simulation models have the benefit of retaining the dependency of the system and any disturbances on time. Algorithms such as Monte Carlo simulation may then be used to explore the variations in the system due to any stochastic properties such as random events or disturbances. However, there is the danger that, even after countless iterations, the system has not been exhaustively explored [2]. This differs from formal verification where the analysis is exhaustive, investigating all possible executions of the system, and therefore can find 'corner cases' which can be missed with alternative approaches. This is particularly important for autonomous vehicles where guarantees on safety are required. In addition, simulation models are time-consuming to build and run. This is especially so when employed in a Monte Carlo simulation of potentially thousands of iterations.

It is therefore of benefit to re-examine the capabilities of formal verification in ensuring that an autonomous vehicle matches its specification. More specifically, we may consider the use of model checking as a verification method. Model checking involves the construction of a mathematical model usually using a formally specified modelling language that captures the system's behaviour and, to verify that the requirements, usually specified in a temporal logic, are satisfied there is a systematic exploration and analysis of the system model. The drawback of using model checking is the state-space explosion problem, and to perform model checking of complex systems abstraction is often required. PRISM [3] is a widely used probabilistic model checker which has support for analysing a wide range of quantitative requirements against a number of different quantitative models. PRISM has been applied to the quantitative verification in a wide range of application domains, from wireless communication protocols to power management systems to systems biology.

An autonomous vehicle may be considered a probabilistic system, with the occurrence of random events determined by some probability. Model checking may then explore the impact of such events on the vehicle mission. That is, the paths the vehicle may take to reach a specific finite state, or the ultimate outcome of the mission. However, the symbolic models employed in model checking software are, by their nature, highly abstracted. Some nuances of the system may be lost in abstraction, where a continuous-time simulation model may be able to retain these nuances due to its greater fidelity.

With quantitative model checking of autonomous vehicles being relatively unexplored territory, we wish to determine the differences between the approaches of formal verification and simulation-based verification. To this end, we employ a simple autonomous system, undergoing a basic mission in a controlled environment. We describe the system and environment by a continuous-time model. We then investigate the impact of probabilistic events by evaluating the results of Monte Carlo simulation. It may be noted that more sophisticated simulation techniques are available, for example, guided M-C simulation for counter-example generation [4, 5, 6]. This method involves attempting to falsify a temporal property by guiding simulation using robustness metrics. Our goal is different. We do not wish to optimise the simulation process itself: the ultimate goal is to produce realistic finite state models for probabilistic verification.

The analysis for this continuous time model are discussed in relation to a concurrent work on a framework for investigate decision-making in autonomous systems, using probabilistic model checking of an abstract model where quantitative data for abstract actions is derived from small-scale simulation models [3]. More precisely, in this work we have built a abstract model of the same autonomous system using this framework. Here we demonstrate that the Monte Carlo simulation results of the detailed continuous-time model for properties, including the expected time mission completion and the probability of certain mission failures, agree with the bounds obtained using the probabilistic model checker PRISM for the abstract discrete-time model.

The paper is structured as follows. Section 2 reviews literature relating to formal verification of autonomous vehicles. Section 3 describes a case scenario, involving an autonomous quadrotor unmanned aerial vehicle (UAV) with a simple search and retrieve mission. Section 4 describes the continuous-time model of the complete system, including the quadrotor, the environment, the target objects to be retrieved. Section 5 describes the guidance system of the quadrotor, which comprises a flight control system (FCS) and finite state machine (FSM), the latter of which provides the vehicle its autonomy. Section 6 presents the results of a single simulation run and those of the Monte Carlo simulation. Section 7 describes how probabilistic model checking may be applied to a highly-abstracted model of our autonomous mission and compares the results obtained through Monte Carlo simulation and probabilistic model checking. Finally, Section 8 presents the conclusions of the research.

2. Related Work

A related area of research is that of reachability analysis for the generation of safe plans for automated vehicles [7, 8, 9]. Dynamic models of automated vehicles are used to generate a safe set of reachable states (with respect to uncertainties in the movement of vehicles in the system), and thence to generate a safe plan of movement. Our goal is different: we start with a dynamic model of a system (a unmanned aircraft), and use it to (ultimately) derive a faithful finite state probabilistic representation of our system. We then use model checking to examine probabilistic properties of the system (by exploring the smaller, finite state model). Not only can we perform reachability analysis using model checking, but we can analyse our system with respect to a wide range of properties expressed via stochastic logic.

Unmanned vehicles are a complex combination of subsystems in both hardware and software domains. As such, model checking may be used in a variety of investigations in unmanned vehicle verification. Typically, the control system of the vehicle is verified such that its effect on the vehicle behaviour meets the desired specification. However, the purpose of this control system can vary between simple automatic stabilisation and complex decision-making algorithms with multiple subsystems. Here, we discuss existing work in formal verification of control systems for vehicles both manned and unmanned.

Examples of the use of model checking for automatic control systems are few. This is likely due to the relative simplicity of such a system in comparison to an autonomous or semi-autonomous system. Typically, control systems such as those on aircraft were easily described as simple autopilots. Many more recent flight control systems (FCS) employ autonomy to some degree and, as such, merit the need for verification. A study by Airbus [10] investigated the use of model checking for a variety of experiments in verification of embedded systems on A340 and A400M aircraft.

A study by Lerda et al. [2] uses model checking to verify of an embedded control system for a UAV. This investigation uses a continuous-time model, developed in MATLAB, to accurately describe the dynamics of the UAV. Model checking is then used to verify the control system. The study highlights the benefits of both simulation and model checking. Simulation can be cumbersome and is often prohibitively time-consuming when exploring the many aspects of a system's behaviour. Conversely, model checking of a continuous-time system is difficult, as the state-space of a continuous system is infinite in size. The use of hybrid automata is one solution to this problem [11, 12, 13], however the formal verification techniques which are applicable to such a model are computationally expensive. Additionally, the complexity of systems which can be analysed using this approach is low. This precludes application to autonomous systems, which are complex by their very nature.

Many instances of model checking for autonomous vehicles involve swarms. Typically, the use of a group of vehicles to perform some task is investigated. Model checking is then used to identify the routes that may be taken to accomplish this task. Examples include UAVs in a pursuer-evader [14] and foraging [15, 16, 17] scenarios. Many of these investigations describe the behaviour of each entity in the swarm as a finite state machine (FSM). In a simple model, the number of entities and their states determine the size of the model's state-space. The dynamics of the vehicles are considered in some cases through the use of simulation [16, 17]. However, in these instances, the vehicles in question are relatively simple in their mechanics.

Other applications of model checking to autonomous vehicles focus on verifying the interaction between agents in co-operative missions [18]. More recent work focusses on the verification of autonomous behaviours with a view to obtaining certification of UAVs for widespread commercial use [1, 19].

3. Scenario

We employ a simple scenario as a case study. A quadrotor UAV is in operation inside a small, constrained environment. This environment is based on the University of Glasgow's Micro Air Systems Technologies (MAST) Laboratory, a $4 \times 7 \times 3$ m cuboidal flight space with an Optitrack[1] motion capture system for tracking UAVs. Three inanimate target objects, each of unique colour and shape, are located somewhere on the floor of the lab. The UAV is programmed such that it should take off from a specified landing site and following a predetermined series of waypoints, while searching for these objects. Upon finding a target, the quadrotor retrieves it from the floor, transports it to a predetermined drop site and deposits it. It then continues following the waypoints until all of the targets have been relocated or it reaches

[1] Natural Point, Inc. https://www.optitrack.com

the final waypoint. It then returns to base, having either succeeded in its mission by relocating all targets or failed by having missed at least one.

Stochastic behaviour is introduced into the scenario through the inclusion of faults, the occurrence of which are based on a uniform probability distribution. Certain system properties are defined similarly, such as the initial locations of the targets, the landing site of the quadrotor and the drop site. These are described in greater detail in Section 4.8.

These random properties can propagate through the model, ultimately affecting the outcome of the mission. The success of the mission is thus a function of the target, drop site and landing site locations and the occurrence of faults. This function is then defined by the model of the autonomous system. This model is a complex combination of hardware and software subsystems.

4. Mathematical Model

The mathematical model describes a multi-agent system and is implemented using an object-oriented programming (OOP) approach. The quadrotor may be considered as either an *active* or *cognitive* agent [20]. The targets can be considered as *passive* agents [20], in that they have their own dynamics but no goals or reactive behaviours. They may alternatively be considered *objects* [21], in that they may be modified by the UAV agent. We choose the former definition for the targets, such that we may denote a generic agent A and use it to specify properties or behaviours common to both UAV and targets. All agents exist with an environment, which corresponds to the MAST Laboratory in this instance.

The UAV agent has more complex behaviours than the targets. Its guidance software contains both a flight control system and decision-making algorithms. It is these decision-making algorithms which allow the UAV to perform its mission with no interference from a human operator.

In this section, we describe the mathematical model of the multi-agent system, comprising the UAV, targets and environment. First, we define the overall structure of the model. That is, how the agents interact with one another and the environment. Then, we define the frames of reference used to describe the behaviours of the agents in a local body-fixed axes system. We then provide the framework for the geometrical models of the simulation, including the environment. Models of the quadrotor hardware are defined, including the vehicle dynamics and sensor systems. Target dynamics are described similarly. Finally, the stochastic variables used to introduce probabilities into the model are detailed.

4.1. Structure

Figure 1 shows a simplified diagram of the multi-agent system. The UAV's sensor suite senses both the internal states of the quadrotor and the geometry of the targets and environment. The sensor measurements are utilised by the guidance system of the UAV, which then provides control inputs to the UAV's rotors. The guidance system comprises several subsystems which allow the UAV to interpret the sensor feedback, decide on a course of action and drive the vehicle towards that action. This is described in greater detail in Section 5. When the UAV grasps a target, the target behaviour then becomes dependent on the response of the UAV. The dynamics of the UAV hence directly impact the target response during grasping. The environment is static and hence unaffected by the behaviour of the agents.

4.2. Frames of Reference

The positions and orientations of agents are described in an inertially-fixed *World* frame of reference, denoted $\mathcal{W}$. This frame is defined such that $z^{\mathcal{W}}$ is in the direction of the local gravity vector, while $x^{\mathcal{W}}$ and $y^{\mathcal{W}}$ are parallel to the boundaries of the environment and satisfy the rule $\hat{\mathbf{x}} = \hat{\mathbf{y}} \times \hat{\mathbf{z}}$. The position of an agent A is described in $\mathcal{W}$ by $\mathbf{r}_{\mathrm{A}} \in \mathbb{R}^3$, where it is assumed that the agent body is rigid.

We can define a second frame of reference $\mathcal{A}$, which is fixed on any agent A and has origin at $\mathbf{r}_{\mathrm{A}}$, which corresponds to the agent centre of mass. The transformation from $\mathcal{W}$ to $\mathcal{A}$ is described by the direction cosine matrix

$$\mathbf{R}_{\mathcal{W}}^{\mathcal{A}} = \begin{bmatrix} \cos\theta\cos\psi & \cos\theta\sin\psi & -\sin\theta \\ \sin\phi\sin\theta\cos\psi - \cos\phi\sin\psi & \sin\phi s_\theta\sin\psi + \cos\phi\cos\psi & \sin\phi\cos\theta \\ \cos\phi\sin\theta\cos\psi + \sin\phi\sin\psi & \cos\phi\sin\theta\sin\psi - \sin\phi\cos\psi & \cos\phi\cos\theta \end{bmatrix} \tag{1}$$

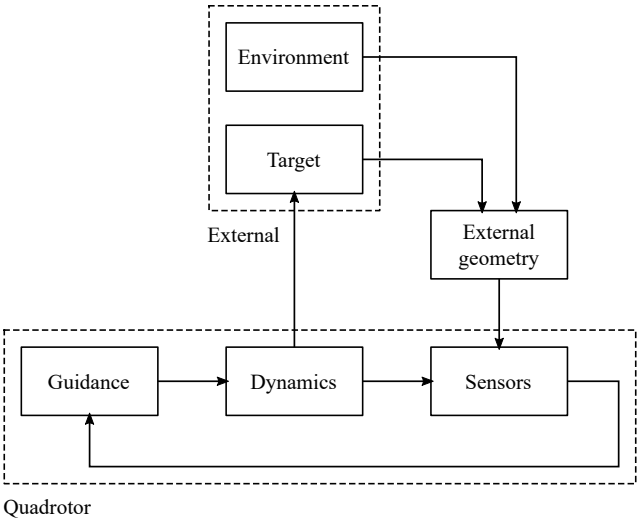

Figure 1: Block diagram of multi-agent system.

Thus, any position $\mathbf{q}^{\mathcal{A}}$ in $\mathcal{A}$ may be described in $\mathcal{W}$ by evaluating the relationship $\mathbf{q}^{\mathcal{W}} = \mathbf{R}^{\mathcal{W}}_{\mathcal{A}}\mathbf{q}^{\mathcal{A}}$. The reverse transformation $\mathbf{q}^{\mathcal{A}} = \mathbf{R}^{\mathcal{A}}_{\mathcal{W}}\mathbf{q}^{\mathcal{W}}$, where $\mathbf{R}^{\mathcal{W}}_{\mathcal{A}} = (\mathbf{R}^{\mathcal{A}}_{\mathcal{W}})^T$ similarly applies.

The orientation is alternatively described by the vector $\boldsymbol{\eta}_{\mathrm{A}} = [\phi, \theta, \psi]^T$. The rotation of $\mathcal{A}$ is described by the angular velocity vector $\boldsymbol{\omega}_{\mathrm{A}} \in \mathbb{R}^3$. The rate of change of orientation $\dot{\boldsymbol{\eta}}_{\mathrm{A}}$ with respect to $\mathcal{W}$ is then described by the relationship

$$\begin{aligned} \dot{\boldsymbol{\eta}}_{\mathrm{A}} &= \mathbf{J}^{-1}_{\mathcal{A}}\boldsymbol{\omega}_{\mathrm{A}} \\ \text{where} \quad \mathbf{J}^{-1}_{\mathcal{A}} &= \begin{bmatrix} 1 & \sin\phi\tan\theta & \cos\phi\tan\theta \\ 0 & \cos\phi & -\sin\phi \\ 0 & \sin\phi\sec\theta & \cos\phi\sec\theta \end{bmatrix} \end{aligned} \tag{2}$$

We can then state that any rigid body agent A in $\mathcal{W}$ has position $\mathbf{r}_{\mathrm{A}}(t) \in \mathbb{R}^3$ and orientation (or attitude) $\boldsymbol{\eta}_{\mathrm{A}}(t) \in \mathbb{R}^3$. The position and attitude responses of a given agent are determined by the dynamics of that agent.

4.3. Geometry

The geometry of the model is defined as a series of polygons with associated faces, vertices and colours. In general, any geometrical object is composed of a series of faces. Each face has an associated colour. The shape, position and orientation of a face in a given reference frame is determined by the number and positions of the vertices associated with that face. The primary use of the geometrical models is to visualise the simulation through animation. However, the models are also used by the UAV's camera sensor model to sense the environment and target agents.

The geometry of any agent A is composed of m vertices, where each vertex $i \in \{1, 2, \ldots m\}$, with position $\mathbf{v}_i \in \mathbb{R}^3$, is associated with one or more faces. If the geometry has n faces, a single face $j \in \{1, 2, \ldots n\}$ of p vertices is then defined by the vector $\mathbf{f}_j$. Each face j has an associated colour, defined by the vector $\mathbf{C}_j \in [0, 1]^3$.

Colour detection is used to identify targets. If we state that each colour $\mathbf{C}_j$ corresponds to a specific face j, we may obtain the indices i of the vertices associated with that face and colour from the elements of $\mathbf{f}_j$. We may then ignore vertices associated with colours outside of a chosen range. The visibility of remaining vertices to the camera sensor is evaluated by considering their positions $\mathbf{v}_i$ in camera image space. This is described in greater detail in Section 5.3.5.

The position $\mathbf{v}^{\mathcal{A}}_i$ of any agent vertex v_i is fixed in the local frame of reference $\mathcal{A}$. It may be described in $\mathcal{W}$ by considering the transformation

$$\mathbf{v}^{\mathcal{W}}_i = \mathbf{R}^{\mathcal{W}}_{\mathcal{A}}\mathbf{v}^{\mathcal{A}}_i + \mathbf{r}_{\mathrm{A}} \tag{3}$$

where $\mathbf{r}_{\mathrm{A}}$ is the position of the origin of $\mathcal{A}$ in $\mathcal{W}$.

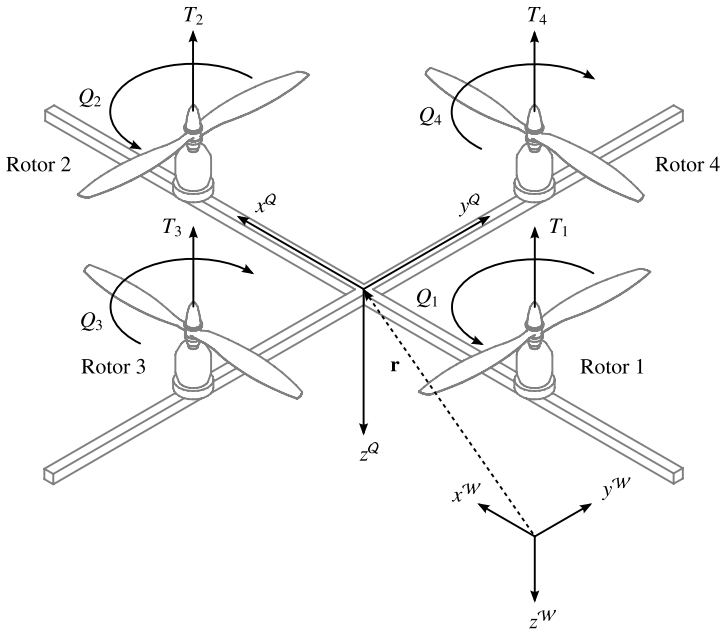

Figure 2: Quadrotor body-fixed frame definition and rotor arrangement.

4.4. Environment

The environment is defined as a simple arena which imposes limits on the freedom of the agents. The World frame $\mathcal{W}$ is defined as beng at the geometric centre of the floor of the environment. The $z^{\mathcal{W}}$-axis is defined in the direction of the local gravity vector. The position of any object or vertex in $\mathcal{W}$ can be described by $\mathbf{r} = [x, y, z]^T$. The environment then imposes the limits $x \in [-2, 2]$, $y \in [-3.5, 3.5]$, $z \in [-3, 0]$.

The geometry of the environment is limited to defining the location of the *drop site* $\mathbf{r}_{\text{ds}}$, where the collected targets are deposited. We may define the *drop zone* as the interior of a circle centred at $\mathbf{r}_{\text{ds}}$ with radius R_{ds}. The UAV perceives targets inside the drop zone, but otherwise ignores them.

4.5. Quadrotor Dynamics

A quadrotor is used as the UAV platform, owing to its versatility and mechanical simplicity. A quadrotor typically consists of four fixed-pitch rotors spaced equally around the centre of mass of the vehicle airframe. The rotors produce both lifting force and, by creating thrust differentials between opposing pairs of rotors, enable control over vehicle attitude. The behaviour of the quadrotor is detailed in [22, 23, 24].

For the search and retrieve scenario, we employ a standard quadrotor UAV with some additional components. The quadrotor is considered a rigid body, with the primary force and moment contribution from the four rotors and the vehicle weight. To facilitate its interaction with any target, we describe the dynamics of a rigidly-mounted grasping arm separately. We then introduce a simple gimbal-stabilised platform for the UAV's camera sensor, and a simple power consumption model.

4.5.1. Rigid-Body Dynamics

We define a quadrotor body-fixed frame $\mathcal{Q}$ which has origin at the vehicle centre of mass, $x^{\mathcal{Q}}$ in the nominal forward direction, $y^{\mathcal{Q}}$ in the nominal starboard direction and $z^{\mathcal{Q}}$ such that the condition $\hat{\mathbf{z}} = \hat{\mathbf{x}} \times \hat{\mathbf{y}}$ is satisfied. The position of the vehicle in $\mathcal{W}$ is denoted $\mathbf{r}_{\text{Q}}$ and its attitude by $\boldsymbol{\eta}_{\text{Q}}$.

Next, we employ a standard quadrotor dynamic model, which assumes a rigid-body response with six degrees of freedom [23, 25]. The rotor dynamics are assumed to be sufficiently faster than the closed-loop body dynamics [23], such that they may be considered instantaneous. A linear relationship between thrust and input PWM signal is assumed [24]. The rotor arrangement shown in Figure 2 is employed.

We make the addition of a simple mass-spring-damper model to represent the vehicle interaction with the floor of the environment. Any collision with the floor is thus described as a highly-damped spring. This "floor force" applies only when the floor plane intersects the rigid geometry of the UAV. The geometry of the UAV is thus characterised by an *effective radius* R_Q around the centre of mass. When the floor plane intersects with this sphere of radius R_Q about $\mathbf{r}_\mathrm{Q} = [x, y, z]^T$, a force is exerted. Otherwise, no force is exerted.

This force, denoted $\mathbf{F}_f$, is assumed to produce no moment and act normal to the floor plane only, that is in the $z^{\mathcal{W}}$ direction. We thus describe it by the conditional relationship

$$\begin{aligned} \mathbf{F}_f &= F_f \mathbf{R}^{\mathcal{Q}}_{\mathcal{W}} \hat{\mathbf{z}} \\ \text{where } F_f &= \begin{cases} -k_\mathrm{Q}(R_\mathrm{Q} + z_\mathrm{Q}) - c_\mathrm{Q}\dot{z} & \text{if } z_\mathrm{Q} \geq -R_\mathrm{Q} \\ 0 & \text{if } z_\mathrm{Q} < -R_\mathrm{Q} \end{cases} \end{aligned} \tag{4}$$

where k_Q is the spring stiffness and c_Q the damping coefficient. These parameters determine the "spring" of the floor. That is, how much the UAV bounces upon collision with the floor.

The quadrotor dynamics are thus described by the 12-state non-linear model

$$\begin{aligned} \ddot{\mathbf{r}}_\mathrm{Q} &= \left(g + \frac{F_f}{m_\mathrm{Q}}\right)\hat{\mathbf{z}} - \frac{K_T}{m_\mathrm{Q}} \sum_{k=1}^{4} u_k \, \mathbf{R}^{\mathcal{W}}_{\mathcal{Q}} \hat{\mathbf{z}} \\ \dot{\boldsymbol{\eta}}_\mathrm{Q} &= \mathbf{J}_{\mathcal{Q}}^{-1} \boldsymbol{\omega}_\mathrm{Q} \\ \dot{\boldsymbol{\omega}}_\mathrm{Q} &= \mathbf{I}_\mathrm{Q}^{-1} \left(\begin{bmatrix} LK_T(u_2 - u_1) \\ LK_T(u_3 - u_4) \\ K_Q(-u_1 - u_2 + u_3 + u_4) \end{bmatrix} - \boldsymbol{\omega}_\mathrm{Q} \times \mathbf{I}_\mathrm{Q} \boldsymbol{\omega}_\mathrm{Q} \right) \end{aligned} \tag{5}$$

where m_Q is the quadrotor mass and $\mathbf{I}$ its inertia tensor, g is the acceleration due to gravity, K_T and K_Q are rotor gains, L is the rotor moment arm and F_f is subject to the conditions imposed by Equation (4). The properties of the quadrotor model are detailed in Table 3. These are obtained from system identification of the Quanser Qball-X4 quadrotor, as described in [24].

4.5.2. Grasper Arm Dynamics

When a target is tethered to the quadrotor, its dynamics are coupled to those of the grasper arm. The dynamics of the arm are thus considered separately from the UAV's rigid-body response. The grasper has fixed position $\mathbf{r}^{\mathcal{Q}}_{\mathrm{G/Q}} \in \mathbb{R}^3$ in the rotating frame $\mathcal{Q}$. It therefore has the inertial position

$$\mathbf{r}_\mathrm{G} = \mathbf{r}_\mathrm{Q} + \mathbf{R}^{\mathcal{W}}_{\mathcal{Q}} \mathbf{r}^{\mathcal{Q}}_{\mathrm{G/Q}}$$

with the second derivative

$$\ddot{\mathbf{r}}_\mathrm{G} = \ddot{\mathbf{r}}_\mathrm{Q} + \mathbf{R}^{\mathcal{W}}_{\mathcal{Q}} \left[\boldsymbol{\omega}_\mathrm{Q} \times \left(\boldsymbol{\omega}_\mathrm{Q} \times \mathbf{r}^{\mathcal{Q}}_{\mathrm{G/Q}} \right) + \dot{\boldsymbol{\omega}}_\mathrm{Q} \times \mathbf{r}^{\mathcal{Q}}_{\mathrm{G/Q}} \right] \tag{6}$$

where translational and rotational accelerations, $\ddot{\mathbf{r}}_\mathrm{Q}$ and $\dot{\boldsymbol{\omega}}_\mathrm{Q}$ respectively, are given by Equation (5). The rotational dynamics of the grasper are identical to that of the quadrotor.

4.5.3. Gimbal Dynamics

The camera sensor is mounted on a gimbal which stabilises it during rolling and pitching motion. When the UAV is level, the principal axis of the camera is in the direction of $z^{\mathcal{W}}$. We define $x^{\mathcal{C}}$ in this direction, while $y^{\mathcal{C}}$ and $z^{\mathcal{C}}$ are in the directions of the camera image's horizontal and vertical directions, respectively.

When the UAV rolls or pitches the gimbal platform rotates to counter this displacement. The roll ϕ_g and pitch θ_g dynamics of the gimbal platform may be described by the first-order system

$$\begin{bmatrix} \dot{\phi}_g \\ \dot{\theta}_g \end{bmatrix} = -\frac{1}{\tau_g} \begin{bmatrix} \phi + \phi_g \\ \theta + \theta_g \end{bmatrix} \tag{7}$$

where τ_g is the time constant of the gimbal rotational response.

It is assumed that the camera centre is coincident with both the centre of rotation of the gimbal platform and the origin of C. The relationship between C and Q is then described by the direction cosine matrix

$$\begin{aligned}\mathbf{R}_Q^C &= \begin{bmatrix} \sin\theta_g & 0 & \cos\theta_g \\ 0 & 1 & 0 \\ -\cos\theta_g & 0 & \sin\theta_g \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\phi_g & \sin\phi_g \\ 0 & -\sin\phi_g & \cos\phi_g \end{bmatrix} \\ &= \begin{bmatrix} \sin\theta_g & -\sin\phi_g\cos\theta_g & \cos\phi_g\cos\theta_g \\ 0 & \cos\phi_g & \sin\phi_g \\ -\cos\theta_g & -\sin\phi_g\sin\theta_g & \cos\phi_g\sin\theta_g \end{bmatrix} \\ \mathbf{R}_C^Q &= \left(\mathbf{R}_Q^C\right)^T \end{aligned} \tag{8}$$

4.5.4. Battery Model

A battery model is employed to allow us to introduce the complication of the UAV periodically returning to base to recharge. As the battery discharges, the voltage supplied to the UAV reduces. This reduction is described by the first-order model

$$\dot{V} = \begin{cases} v_c & \text{if charging and } V < V_{\max} \\ v_d & \text{if not charging} \\ 0 & \text{if } V \geq V_{\max} \end{cases} \tag{9}$$

where $v_c > 0$ is the charging rate, which applies when the UAV is idle at the landing site, $v_d < 0$ is the discharge rate, which applies when the UAV is pursuing its mission. The voltage level has the upper limit $V \leq V_{\max}$, above which it will not charge. The discharge rate is approximated as a constant. In reality, the battery drain is affected by the power requirements of the rotors and other hardware components.

4.5.5. Stochastic Properties of the Quadrotor

We have previously described a simple deterministic quadrotor model. The model has enough detail that it may realistically describe a quadrotor which is capable of sitting idle on the ground, flying and grasping objects. In reality, the quadrotor is subject to a number of stochastic events. Through their effect on the system behaviour, these events can impact the particulars of the mission, including its outcome. Here, we discuss some of these events and, if relevant, how we have chosen to incorporate these into our model.

First, we may consider the effect of external disturbances. Such disturbances are products of the local environment and are governed by a system far greater in complexity than the quadrotor itself. Thus, we may consider them to be random. These disturbances may manifest as forces or moments acting on the vehicle body. They may also impact the performance of the rotors by changing the local airflow around the rotor disk. A robust control system is typically employed to reduce the impact of such disturbances. In reality, even with robust control, these disturbances may have the effect of changing the path the UAV takes, or the time it takes to perform a manoeuvre.

The quadrotor system may also be subject to internal disturbances. These may manifest as a change in a parameter which is expected to be constant. For example, it is assumed that the rotor thrust gain K_T does not change. However, the thrust gain can be altered by a number of phenomena, including varying battery level and local airflow around the rotor, as indicated above. One or more rotors may also fail entirely, due to a fault in any of the components between autopilot and rotor, such as the motor or propeller. Regardless of where the fault occurs, the effect is a complete loss of thrust and torque from that rotor.

The occurrence such an *actuator fault* is based on a large number of variables. We abstract the probability of such a fault occurring by using a geometric distribution assuming that, with the fixed period T_a, the probability of an actuator fault equals $P_a : \mathbb{R}_{\geq 0} \to [0, 1]$ (see Table 3). We model the actuator fault as a complete loss of thrust in a single rotor, that is $K_{Tk} = 0$ for a random rotor $k \in \{1, 2, 3, 4\}$. For example, if a fault occurs in the rear rotor $k = 1$, the maximum propulsive force is reduced by a quarter and the u_1 terms in Equation (5) may be neglected.

While control strategies are available to deal with an actuator fault [26], the reduced capability of the vehicle typically results in a controlled emergency landing, if not a crash landing. In Section 5.3, we introduce a control algorithm specifically for emergency landings after actuator loss.

A fault in the UAV's grasping mechanism may result in the mechanism becoming stuck in its current state or it may release any tethered objects. Since the former may result in a target being permanently tethered to the UAV and thus preventing mission success, we limit our model to include the later case only. It is assumed that some fault occurs which causes the grasper mechanism to deactivate. This has no effect when no target is being grasped. However, if a target is held by the mechanism, it is released and the UAV must retrieve it again to continue the mission. Just as we define the probability of an actuator fault occurring, we may assume that the probability of a *grasper fault* occurring within a fixed time period T_g equals $P_g : \mathbb{R}_{\geq 0} \to [0, 1]$. The impact on the UAV behaviour due to a grasper fault is detailed in Section 5.1.

Finally, we randomly define the pose (position and heading) of the UAV at the beginning of the mission. This has the effect of altering the behaviour of the UAV as the predicates which trigger transitions from mode to another may occur at different times and under different circumstances. That the target locations and drop site are also randomly defined further contributes to the uncertainty in the system.

4.6. Quadrotor Sensors

The quadrotor employs a standard sensor suite to measure system states and sense the environment. A motion capture system allows direct measurement of vehicle position and attitude. An inertial measurement unit (IMU), containing tri-axial accelerometers and gyroscopes, measures accelerations and angular rates in $\mathcal{Q}$. A camera sensor mounted on the underside of the vehicle is used to visually sense the geometry of the environment and targets. This section describes the models of the sensor behaviour. Interpretation and exploitation of sensor outputs is described in Section 5.

4.6.1. Motion Capture System

The position and attitude of the UAV in $\mathcal{W}$ are measured by the motion capture system. The model is based on the MAST Laboratory's Optitrack system . Neglecting measurement noise, we may assume the motion capture system measures both position and attitude with zero error, giving

$$\hat{\mathbf{o}}_r = \mathbf{r}_{\mathrm{Q}}, \qquad \hat{\mathbf{o}}_\eta = \boldsymbol{\eta}_{\mathrm{Q}} \tag{10}$$

4.6.2. Inertial Measurement Unit

The angular rates of the UAV in $\mathcal{Q}$ are measured by the onboard IMU. Again neglecting any noise or bias errors, we describe the gyroscope measurement by

$$\hat{\mathbf{g}} = \boldsymbol{\omega}_{\mathrm{Q}} \tag{11}$$

4.6.3. Camera Sensor

The camera sensor provides an image of any unobscured object within its field of view. In reality, the camera captures an image which is then used in an object detection algorithm. Examples of such algorithms range from simple colour or intensity detection to more complex shape or pattern recognition. The camera model is described in such a way that it enables the process of colour detection to be emulated. This process is further detailed in Section 5.3.5.

As detailed in Section 4.3, any vertex i of a geometrical model has position $\mathbf{v}_i^{\mathcal{W}}$ in $\mathcal{W}$. The camera operates by projecting geometry within the its field of view onto an *image plane*, demonstrated in Figure 3. In order to achieve this projection, the vertex position must be determined in a frame $\mathcal{C}$, fixed on the camera body.

The camera has fixed position $\mathbf{r}^{\mathcal{Q}}_{\mathrm{C/Q}}$ and orientation $\mathbf{R}^{\mathcal{Q}}_{\mathcal{C}}$ in $\mathcal{Q}$. Any vertex $\mathbf{v}^{\mathcal{W}}$ in $\mathcal{W}$ thus has position, relative to the origin of $\mathcal{C}$, of

$$\mathbf{v}^{\mathcal{C}} = \mathbf{R}^{\mathcal{C}}_{\mathcal{Q}} \left[\mathbf{R}^{\mathcal{Q}}_{\mathcal{W}} (\mathbf{v} - \mathbf{r}_{\mathcal{Q}}) + \mathbf{r}^{\mathcal{Q}}_{\mathrm{C/Q}} \right] \tag{12}$$

The vertex is now defined relative to the camera and centre. If $\mathbf{v}^{\mathcal{C}} = [v_x, v_y, v_z]^T$, then the projection of $\mathbf{v}^{\mathcal{C}}$ onto the image plane is described by the coordinates

$$x_c = f\frac{v_y}{v_x}, \qquad y_c = -f\frac{v_z}{v_x} \tag{13}$$

where f is the focal length of the camera.

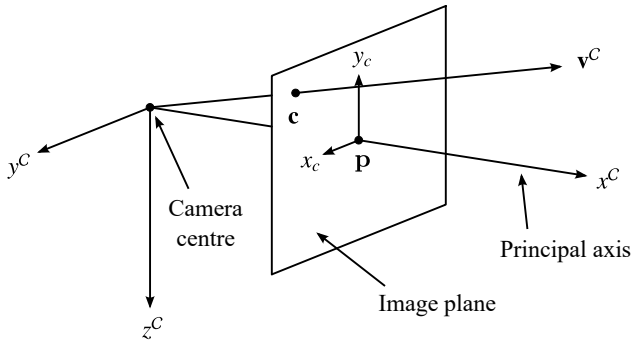

Figure 3: Geometry of a pinhole camera. The coordinates of a point with position $\mathbf{r}$ in Euclidean 3-space are mapped to 2-space be considering the interaction of the point with the image plane [27].

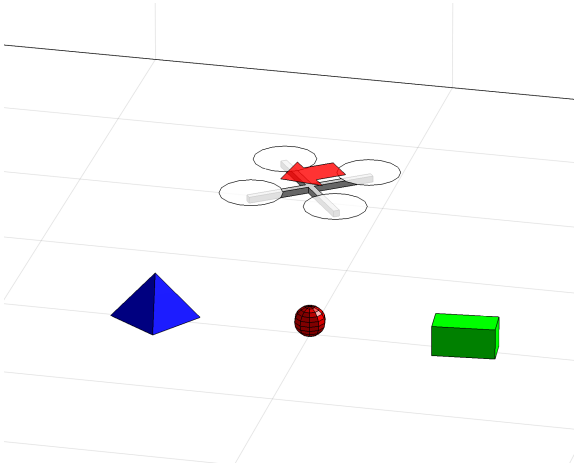

Figure 4: Target geometries and colours as rendered by simulation visualisation.

The camera has horizontal field of view λ and aspect ratio A. For any given vertex $\mathbf{c} = [x_c, y_c]^T$ to be visible to the camera sensor, it must be within the field of view. The conditions

$$\begin{aligned} -f \tan\lambda \leq x_c \leq f \tan\lambda \\ -\frac{f}{A} \tan\lambda \leq y_c \leq \frac{f}{A} \tan\lambda \end{aligned} \tag{14}$$

must therefore be satisfied.

4.7. Target Model

While static for the majority of the mission, each target is given dynamic behaviours in order to facilitate its interaction with the UAV. In isolation, the behaviour of a target is purely deterministic. However, any interactions with the UAV are subject to the influence of stochastic and probabilistic behaviours.

While each target has the same dynamic behaviour, we consider three distinct shapes and colours. This allows the UAV to discriminate between targets based on their colour and, in the event of more complex object detection being required, their shape. The targets may be seen in Figure 4. We model the three targets as a blue pyramid, red sphere and green cuboid, respectively.

4.7.1. Rigid-Body Dynamics

We may consider each target to be a rigid body of mass m_T. A target body-fixed frame $\mathcal{T}$ is defined with origin at the centre of mass of the target and $\mathbf{R}_{\mathcal{T}}^{\mathcal{W}} = \mathbf{I}_3$ before grasping occurs, where $\mathbf{I}_3$ is the 3×3 identity matrix.

When tethered to the UAV during grasping, the dynamics of the target are identical to those of the UAV grasping arm, and are therefore given by

$$\begin{aligned} \ddot{\mathbf{r}}_{\mathrm{T}} &= \ddot{\mathbf{r}}_{\mathrm{G}} \\ \dot{\boldsymbol{\eta}}_{\mathrm{T}} &= \mathbf{J}_{\mathcal{T}}^{-1}\boldsymbol{\omega}_{\mathrm{T}} \\ \dot{\boldsymbol{\omega}}_{\mathrm{T}} &= \dot{\boldsymbol{\omega}}_{\mathrm{Q}} \end{aligned} \tag{15}$$

where $\ddot{\mathbf{r}}_{\mathrm{G}}$ and $\dot{\boldsymbol{\omega}}_{\mathrm{Q}}$ are defined by Equations (6) and (5), respectively.

When not tethered to the UAV, the target dynamics are driven by a simple gravitational force and floor model

$$\begin{aligned} \ddot{\mathbf{r}}_{\mathrm{T}} &= g\hat{\mathbf{z}} + \frac{1}{m_{\mathrm{T}}}\mathbf{F}_f \\ \dot{\boldsymbol{\eta}}_{\mathrm{T}} &= \mathbf{J}_{\mathcal{T}}^{-1}\boldsymbol{\omega}_{\mathrm{T}} \\ \dot{\boldsymbol{\omega}}_{\mathrm{T}} &= \mathbf{0} \end{aligned} \tag{16}$$

where $\mathbf{F}_f$ is the force exerted by the floor on the target. This now includes a frictional component in the horizontal plane and is given by

$$\mathbf{F}_f = \begin{cases} \mathbf{R}_{\mathcal{W}}^{\mathcal{T}}\left(-k_{\mathrm{T}}\left(R_{\mathrm{T}} + z_{\mathrm{T}}\right)\hat{\mathbf{z}} - c_{\mathrm{T}}\dot{\mathbf{r}}_{\mathrm{T}}\right) & \text{if } z_{\mathrm{T}} \geq -R_{\mathrm{T}} \\ \mathbf{0}_{3\times 1} & \text{if } z_{\mathrm{T}} < -R_{\mathrm{T}} \end{cases} \tag{17}$$

where R_{T} is the effective radius of the target and k_{T}, c_{T} represent the spring stiffness and damping of the floor-target interaction. This floor model allows us to describe the behaviour of the target as it impacts the floor and comes to rest.

4.7.2. Stochastic Properties of the Target

The internal behaviour of the target as described by Equation (16) is purely deterministic. However, when grasped by the UAV, its behaviour is determined by that of the UAV (Equation (15)). In this state, any stochastic behaviours demonstrated by the UAV will similarly impact the target.

The actuator fault described previously has the effect of prematurely ending the mission, thus its effect on the target is largely inconsequential. Conversely, a fault in the grasping mechanism has significant impact on both the UAV and target.

When a fault occurs in the UAV's grasping mechanism, any target currently tethered to the UAV is released. The target then follows a ballistic trajectory before colliding with the ground. In this phase, the state transition of the target is purely deterministic (Equation 16), although the state vector at the beginning of the trajectory is determined by the state of the UAV at release, which is stochastic. The eventual resting place of the target after being dropped is thus similarly random.

Similarly, the initial pose of each target is randomly defined. This impacts the behaviour of the UAV and thus the time and location at which the target is picked up by the UAV and deposited in the drop zone.

4.8. Summary of Modelled Stochastic Properties

As detailed previously, both the UAV and target agents are subject to non-deterministic behaviours. These may be caused by randomly-define properties or by events which have a certain probability of occurring.

4.8.1. Random Properties

As described in the models for each agent type, the initial positions of the UAV and targets on the $x^{\mathcal{W}}$-$y^{\mathcal{W}}$ plane are defined at random. The headings ψ are similarly specified. If an agent A has position $\mathbf{r}_{\mathrm{A},0} = [x_0, y_0, z_0]^T$ at time $t = 0$, then the position in the horizontal plane is defined randomly in the ranges $x_0 \in [-1.9, 1.9]$ and $y_0 \in [-3.4, 3.4]$. The initial height is defined $z_0 = -R_{\mathrm{A}}$, where R_{A} is the effective radius of the agent. The drop site location $\mathbf{r}_{\mathrm{ds}} = [x_{\mathrm{ds}}, y_{\mathrm{ds}}, z_{\mathrm{ds}}]^T$ is similarly defined, with $x_{\mathrm{ds}} \in [-1.5, 1.5]$, $y_{\mathrm{ds}} \in [-3, 3]$ and $z_{\mathrm{ds}} = 0$. The random properties are generated using a uniform distribution.

4.8.2. Probabilistic Events

We have described the effects of actuator and grasper mechanism faults on the UAV system. We may implement these faults in our simulation by evaluating the probability of their occurrence in a given time step dt.

An actuator fault may occur at any time. If the UAV is airborne, an actuator fault forces the UAV to perform an emergency landing, thus failing the mission. The probability of an actuator fault occurring in some period of time T_a is denoted P_a. The probability of it occurring during any given time step is then defined by

$$P_{a,dt} = 1 - (1 - P_a)^{\frac{dt}{T_a}}$$

A fault in the grasping mechanism may also occur at any time. However, it only affects the mission when a target is tethered to the UAV. Upon this fault occurring, the target is dropped. The probability of a grasping mechanism fault occurring in some period of time T_g is denoted P_g. The probability of it occurring during any given time step is then

$$P_{g,dt} = 1 - (1 - P_g)^{\frac{dt}{T_g}}$$

Finally, we may define a generic *system fault*. This represents the detection of some undesirable condition after the UAV has taken off. Examples of such conditions include interrupted communications between UAV and ground control station, erratic sensor feedback or the absence of some required system parameters or commands. The UAV's decision-making system has an *Initialise* mode specifically for detecting these conditions. The probability that some undesirable condition is detected during this mode is given by $P_s : \mathbb{R}_{\geq 0} \to [0, 1]$ and is independent of time.

5. Autonomous Guidance System

The guidance system of the UAV consists of a number of autonomous *modes* which dictate the behaviour of the UAV at any given time. Each mode comprises a combination of automatic control and trajectory modules, which collaborate to drive the UAV towards a specific position and heading. The modes and the transitions between them may be represented as a finite state machine, shown in Figure 5. The UAV can respond to any internal or external event considered in the simulation. In a real-world scenario, events may occur which are unmodelled in the simulation. In implementing this guidance system in reality, this would make the UAV only semi-autonomous. Human intervention would likely be required to mitigate the effects of some events which the guidance system is not designed to react to.

Using the FSM to represent the behaviour of the guidance system, each mode corresponds to a state of the FSM. Only one mode may thus be active at any given time. The guidance system progresses from one mode to another under certain conditions. The criteria which dictate an exit from the current mode and which mode succeeds the current one are known as *predicates*. The predicates for each mode are shown with reference to the FSM in Figure 5.

We now detail the behaviour of each mode with reference to the control and trajectory modules or commands employed in each case. The algorithms of each of these modules are then described.

5.1. Guidance Modes

The guidance modes determine which control and trajectory modules are active at any given time, and what the inputs to each of these modules are. Each mode thus comprises a unique set of instructions and modules. These are summarised in Table 1.

The UAV begins the mission in *Idle*. It then progresses to *Take-off* and subsequently *Initialise*, where a self-diagnosis is performed. The probability of continuing the mission (as opposed to landing) is determined by the system fault probability P_s. The UAV may then either progress to *Search* or *Land* mode, whereupon it continues as per the predicates of each state. The modes, transitions and predicates are detailed in Figure 5. Each mode may be summarised as follows.

5.1.1. Idle

The UAV sits at rest with its rotors inactive. The UAV begins the mission in this mode, resting at the landing site. It returns to this mode after completing the mission, with one of two landing modes leading to *Idle*. When succeeding *Land*, the UAV rests at the landing site. When succeeding *Emergency land*, its location is wherever it comes to rest after colliding with the floor.

The UAV exists this mode only when $V = V_{\max}$. This ensures that the battery is fully charged before beginning or resuming the mission. It may then progress to *Take-off* if the mission has not yet been completed.

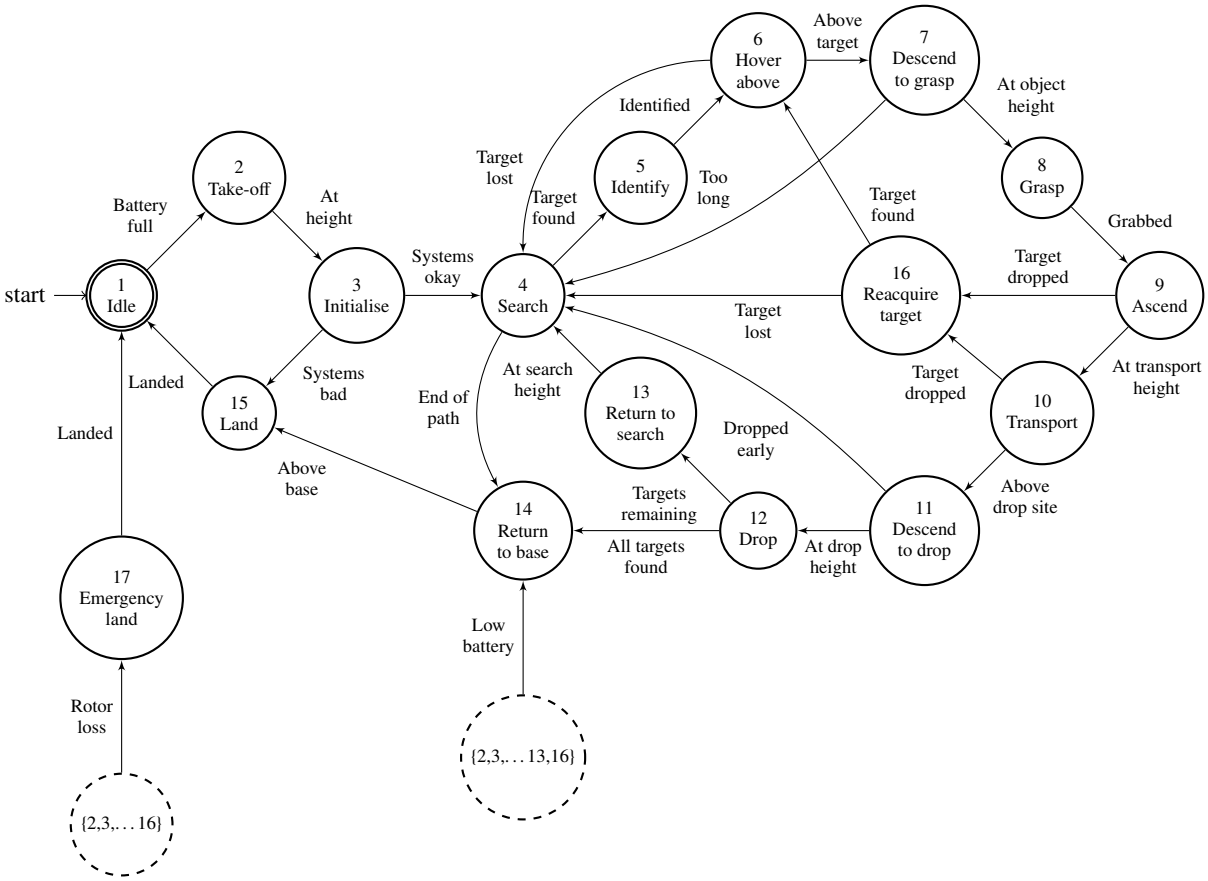

Figure 5: Finite-state machine describing AI logic.

Table 1: Guidance commands and modules for each AI mode.

Mode	Input Command / Module	Trajectory Command / Module
Idle	$\mathbf{u}_{\mathrm{Q}} = [0, 0, 0, 0]^T$	–
Take-off	State feedback controller	$\mathbf{r}_d = [x_0, y_0, z_{\mathrm{hvr}}]^T$
Initialise	State feedback controller	$\mathbf{r}_d = [x_0, y_0, z_{\mathrm{hvr}}]^T$
Search	State feedback controller	Search pattern
Identify	State feedback controller	$\mathbf{r}_d = \mathbf{r}_{\mathrm{entry}}$
Hover above target	State feedback, visual controllers	Object tracking, $z_d = z_{\mathrm{srch}}$
Descend to grasp	State feedback, visual controllers	Object tracking, $z_d = -(z_{\mathrm{G/Q}} + 2R_{\mathrm{T}})$
Grasp	State feedback controller	$\dot{x}_d = 0, \dot{y}_d = 0, z_d = -(z_{\mathrm{G/Q}} + 2R_{\mathrm{T}})$
Ascend	State feedback controller	$\dot{x}_d = 0, \dot{y}_d = 0, z_d = z_{\mathrm{trnsprt}}$
Transport	State feedback controller	$\mathbf{r}_d = [x_{\mathrm{ds}}, y_{\mathrm{ds}}, z_{\mathrm{trnsprt}}]^T$
Descend to drop	State feedback controller	$\mathbf{r}_d = [x_{\mathrm{ds}}, y_{\mathrm{ds}}, -(z_{\mathrm{G/Q}} + 2R_{\mathrm{T}})]^T$
Drop	State feedback controller	$\mathbf{r}_d = [x_{\mathrm{ds}}, y_{\mathrm{ds}}, -(z_{\mathrm{G/Q}} + 2R_{\mathrm{T}})]^T$
Return to search	State feedback controller	$\mathbf{r}_d = [x_{\mathrm{ds}}, y_{\mathrm{ds}}, z_{\mathrm{srch}}]^T$
Return to base	State feedback controller	$\mathbf{r}_d = [x_0, y_0, z_{\mathrm{hvr}}]^T$
Reacquire target	State feedback controller	$\mathbf{r}_d = \mathbf{r}_{\mathrm{entry}}$
Land	State feedback controller	$\mathbf{r}_d = [x_0, y_0, -R_{\mathrm{Q}}]^T$
Emergency land	Emergency contoller	$z_d = -R_{\mathrm{Q}}$

5.1.2. Take-off

The UAV takes off to a position above the landing site $\mathbf{r}_d = [x_0, y_0, z_{\text{hvr}}]^T$, where z_{hvr} is the hover height. This ensures that the ground effect region is cleared before performing further manoeuvres. It then proceeds to *Initialise* upon satisfying the condition $\|\mathbf{r}_{\text{Q}} - \mathbf{r}_d\| < tol$ where tol is a tolerance value which varies from case to case.

5.1.3. Initialise

The UAV performs a self-diagnosis to identify any system faults. A system fault is detected with probability P_s, in which case the UAV proceeds to *Land*. If no fault is detected, the UAV makes a transition to *Search*.

5.1.4. Search

The UAV follows a series of waypoints as defined by the *Search pattern* trajectory module. The spacing of the waypoints and the search height z_{srch} are defined such that the camera's field of view overlaps as it passes between pairs of waypoints. The velocity command is limited to $v_{\text{max}} = 2$ m/s to ensure that the UAV does not pass over targets too quickly. The *Object tracking* module is active during this mode, but does not inform the controller. It provides only a boolean $D \in \{0, 1\}$ which specifies whether a target has been detected. When $D = 1$, a transition to *Identify* is triggered (see Figure 5).

As the UAV is aware of how many targets are present in the environment, it will return to base if all targets have been located and retrieved. In such a situation, it does not therefore reach the final waypoint. If it reaches the final waypoint, it is assumed that not all targets have been found. The UAV then proceeds to *Return to base* and triggers a *Mission failed* flag. A time limit T_{max} defined for the mission triggers a similar behaviour.

5.1.5. Identify

The transition to *Identify* is triggered when a target is detected by the *Object tracking* module, as described in Section 5.3.5. The UAV typically has some lateral motion during this transition. The controller thus compensates for overshooting the target by returning the UAV to the position recorded as it entered *Identify*, denoted $\mathbf{r}_{\text{entry}}$. Once it has returned within a small radius of this position, that is $\|\mathbf{r}_{\text{Q}} - \mathbf{r}_{\text{entry}}\| < tol$, the UAV progresses to *Hover above target*.

5.1.6. Hover Above Target

As a result of the position command issued during *Identify*, the UAV is now in a position above the target which makes the target visible to the camera sensor (see Figure 6). The *Object tracking* module detects blocks of colour within a certain RGB range and determines the coordinates of their centroids. The coordinates of the centroid closest to the centre of the image are supplied to the *Visual controller*. This has the result of moving the UAV directly above the target. When the vehicle velocity in the $x^{\mathcal{W}}$-$y^{\mathcal{W}}$ plane is near-zero, that is $\sqrt{\dot{x}_{\text{Q}}^2 + \dot{y}_{\text{Q}}^2} < tol$, the UAV is above the target and the UAV proceeds to *Descend to grasp*.

5.1.7. Descend to Grasp

The UAV uses the *Object tracking* and *Visual controller* modules to maintain its position directly above the target, while descending to retrieve it. When the grasping mechanism position is within a small radius of the top of the target, that is $\|\mathbf{r}_{\text{G}} - (\mathbf{r}_{\text{T}} - R_{\text{T}})\| < tol$, the UAV makes a transition to *Grasp*.

5.1.8. Grasp

The grasping mechanism is activated and the dynamics of the grasped target are tethered to those of the UAV, as described by Equation (15). The UAV then proceeds to *Ascend*.

5.1.9. Ascend

The UAV ascends to the transport height z_{trnsprt}, while maintaining zero velocity in the horizontal plane. The UAV proceeds to *Transport* upon reaching z_{trnsprt}, that is, satisfying $|z_{\text{Q}} - z_{\text{trnsprt}}| < tol$. A grasping mechanism fault occurs in this mode with probability P_g. When such a fault is detected, the UAV moves instead to *Reacquire target*.

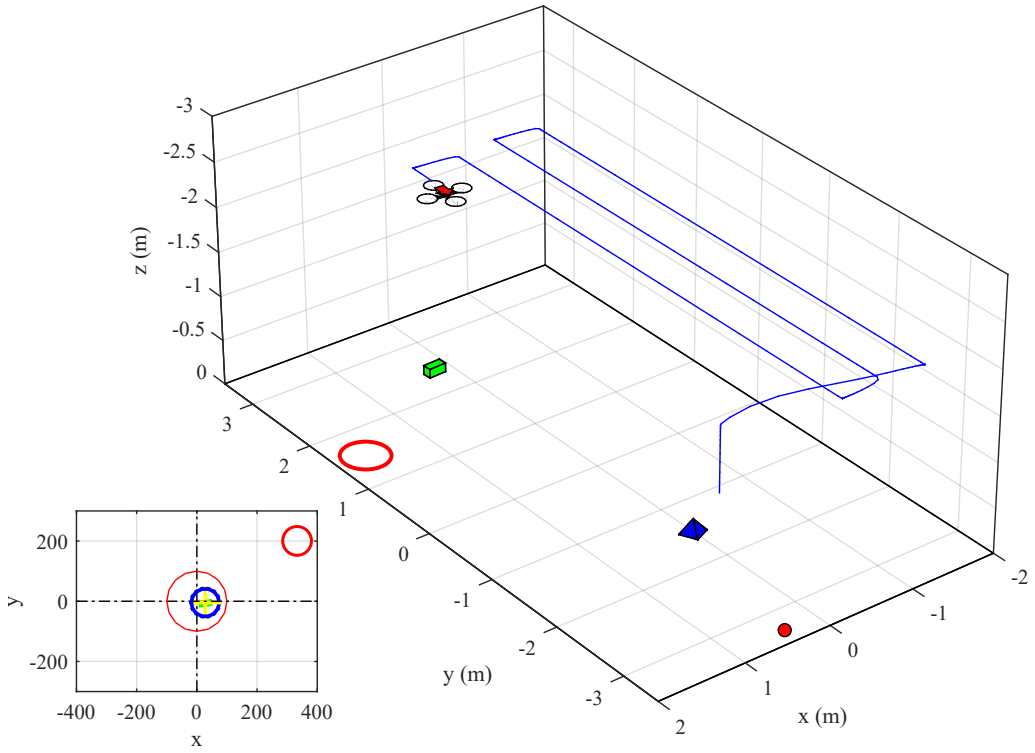

Figure 6: Mission visualisation, showing UAV, targets, UAV flight path and camera sensor image.

5.1.10. Transport

The UAV transports the tethered target to a position above the drop site, defined by $\mathbf{r}_d = [x_{\text{ds}}, y_{\text{ds}}, z_{\text{trnsprt}}]^T$. The UAV proceeds to *Descend to drop* upon the UAV reaching this position, that is $\|\mathbf{r}_{\text{Q}} - \mathbf{r}_d\| < tol$, for the mode-specific command $\mathbf{r}_d$. A grasping mechanism fault occurs in this mode with probability P_g. Again, when such a fault is detected, a transition to *Reacquire target* is triggered.

5.1.11. Descend to Drop

The UAV descends to a height where the target is touching the floor, that is $z_d = -(z_{\text{G/Q}} + 2R_{\text{T}})$. If the target is dropped early during this mode due to a grasping mechanism fault, the target is recorded as having been successfully deposited and the UAV proceeds to *Return to search*. Otherwise, it progresses to *Drop* upon satisfying $\|\mathbf{r}_{\text{Q}} - \mathbf{r}_d\| < tol$, where $\mathbf{r}_d = [x_{\text{ds}}, y_{\text{ds}}, z_d]^T$.

5.1.12. Drop

The grasping mechanism is deactivated and the dynamics of the target are untethered from the UAV and described again by Equation (16). The UAV proceeds to *Return to search* if the mission is not complete (i.e. there are targets still to be found) or to *Return to base* if all targets have been deposited at the drop site.

5.1.13. Return to Search

The UAV ascends vertically to search height above the drop site, that is, the position $\mathbf{r}_d = [x_{\text{ds}}, y_{\text{ds}}, z_{\text{srch}}]^T$. It then proceeds to *Search* upon satisfying $\|\mathbf{r}_{\text{Q}} - \mathbf{r}_d\| < tol$.

5.1.14. Return to Base

The UAV returns to a position $\mathbf{r}_d = [x_0, y_0, z_{\text{hvr}}]^T$ above the landing site upon either successfully depositing all targets, reaching the end of the search path or receiving a low battery warning. A transition to *Land* is triggered upon satisfying $\|\mathbf{r}_{\text{Q}} - \mathbf{r}_d\| < tol$.

5.1.15. Land

The UAV descends from its position above the landing site to rest on the landing site itself, at position $\mathbf{r}_d = [x_0, y_0, -R_Q]^T$. The UAV then returns to *Idle* upon satisfying the conditions $\|\mathbf{r}_Q - \mathbf{r}_d\| < tol_1$ and $\|\dot{\mathbf{r}}_Q\| < tol_2$. This ensures that the motors are set to idle when the UAV is stationary at the landing site.

5.1.16. Reacquire Target

A transition to *Reacquire target* from *Ascend* or *Transport* is triggered upon a grasping mechanism fault occurring. The transition occurs immediately upon the grasping mechanism failing. As the UAV may have a high velocity when the transition occurs, it is commanded to return to the position $\mathbf{r}_{\text{entry}}$ at which it entered the *Reacquire target* mode in order to maximise the probability of visually reacquiring the target. To exit this mode, the conditions $\|\mathbf{r}_Q - \mathbf{r}_{\text{entry}}\| < tol_1$, $\|\dot{\mathbf{r}}_Q\| < tol_2$ must be satisfied. The guidance mode changes to *Hover above target* if the *Object tracking* module detects a target, that is, the module has set $D = 1$, or to *Search* if the previous conditions are satisfied and no target is detected, that is, $D = 0$. That is, the target has been lost and the highest probability of reacquiring it occurs while following the search waypoints.

5.1.17. Emergency Land

A transition to *Emergency land* may occur in any mode bar *Emergency land* itself and *Idle*. The *Emergency controller* is employed, accepting only the height command $z_d = -R_Q$. The UAV returns to *Idle* upon the condition $|z_Q + R_Q| < tol$ being satisfied.

5.2. Additional Transitions

As highlighted by Figure 5, the modes *Reacquire target* and *Emergency land* may be reached from a large number of other modes. These transitions are triggered by the detection of faults. The autonomous controller reacts to the occurrence of faults by switching to a mode which is designed to handle the effects of the fault.

5.2.1. Low Battery Warning

A simple battery monitor subsystem checks the battery voltage V and triggers a transition to *Return to base* when $V \leq V_{\text{th}}$, where V_{th} is the lower threshold. This transition can occur in any mode except those in the path leading from *Return to base* to *Idle*, that is *Return to base*, *Land*, *Emergency land* or *Idle*. Following this transition, any tethered targets are released, as the mission priority is to recharge the battery.

5.2.2. Actuator Fault

An actuator fault occurs with probability P_a. This triggers a transition to *Emergency land*. It may occur in any mode except *Emergency land* or *Idle*.

5.3. Control and Trajectory Modules

The control and trajectory modules active at any given time are determined by the UAV's mode, as previously described. The modules work in combination to follow the commands issued in the active mode. As an example, during *Take-off*, the desired vehicle behaviour is to hover at a static position above the landing site. A constant position command is therefore supplied to the *State feedback controller*. In *Search*, the UAV is required to follow one waypoint after another. It therefore requires a *Search pattern* module to define the changing position commands and the *State feedback controller* to track them. A *State reconstruction* module is active at all times. This module is required to reconstruct the dynamic states of the quadrotor from the available sensor measurements.

The vehicle control laws determine the *pseudo-inputs* of the UAV, $\hat{\mathbf{u}} = [u_{\text{col}}, u_{\text{roll}}, u_{\text{pitch}}, u_{\text{yaw}}]^T$, detailed in [28, 29, 23]. We may define these four pseudo-inputs in terms of the true rotor inputs $\mathbf{u} = [u_1, u_2, u_3, u_4]^T$ such that

$$\hat{\mathbf{u}} = \mathbf{C}\mathbf{u}, \quad \text{where } \mathbf{C} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & -1 \\ -1 & 1 & 0 & 0 \\ -1 & -1 & 0 & 0 \end{bmatrix} \tag{18}$$

If the controller determines the pseudo-inputs, then the true rotor inputs are found from the inverse relationship $\mathbf{u} = \mathbf{C}^{-1}\hat{\mathbf{u}}$.

This section describes the control and trajectory modules which work in combination to supply these pseudo-inputs and therefore drive the UAV behaviour. The module combinations and mode-dependent commands are specified in Table 1.

5.3.1. State Reconstruction

The control and navigation systems of the UAV require knowledge of the system states $\mathbf{x}_\mathrm{Q}$. These are not directly obtainable and must be measured through the sensors. The states may then be reconstructed in software from the sensor outputs. The sensor model omits errors in the measurement relationships. It is therefore assumed that the position $\mathbf{r}_\mathrm{Q}$ and attitude $\boldsymbol{\eta}_\mathrm{Q}$ are perfectly reconstructed from the Optitrack measurements, while the angular rates $\boldsymbol{\omega}_\mathrm{Q}$ are similarly reconstructed without error from the gyroscope outputs. Translational velocity is obtained by filtering and differentiating the vehicle position

$$\hat{\mathbf{v}}_\mathrm{Q} = \frac{Ns}{s+N}\mathbf{r}_\mathrm{Q} \tag{19}$$

where N is the filter bandwidth.

5.3.2. State Feedback Controller

A state feedback controller with feedback linearisation [28, 23, 30] stabilises the quadrotor and ensures accurate tracking of position and yaw commands. The position controller consists of cascaded state feedbacks, with the desired acceleration given by

$$\ddot{\mathbf{r}}_d = K_{dr}\left(\dot{\mathbf{r}}_d - \dot{\mathbf{r}}\right) \tag{20}$$

and the desired velocity given by

$$\dot{\mathbf{r}}_d = K_{pr}\left(\mathbf{r}_d - \mathbf{r}_\mathrm{Q}\right) \tag{21}$$

where $\mathbf{r}_d$ is the desired position and $\dot{\mathbf{r}}_d$ is subject to the limit $\|\dot{\mathbf{r}}_d\| \leq v_\mathrm{max}$, where v_max is an artificial maximum velocity limit. Decoupling of the position and velocity feedbacks is required for the visual controller described by Equation (29). Linearising feedbacks then use the desired acceleration command $\ddot{\mathbf{r}}_d$ to determine the collective pseudo-input

$$u_\mathrm{col} = \frac{m_\mathrm{Q}(g - \ddot{z}_d)}{K_T \cos\phi\cos\theta} \tag{22}$$

and the roll and pitch commands

$$\begin{aligned} \phi_d &= \arcsin\left(\frac{m_\mathrm{Q}\left(\ddot{y}_d\cos\psi - \ddot{x}_d\sin\psi\right)}{K_T u_\mathrm{col}}\right) \\ \theta_d &= -\arcsin\left(\frac{m_\mathrm{Q}\left(\ddot{x}_d\cos\psi + \ddot{y}_d\sin\psi\right)}{K_T u_\mathrm{col}\cos\phi}\right) \end{aligned} \tag{23}$$

where the commands are subject to the limits $\{|\phi_d|, |\theta_d|\} \leq a_\mathrm{max}$, where a_max is an artificial maximum roll/pitch limit. The attitude controller is similarly defined by the state feedback

$$\dot{\boldsymbol{\omega}}_d = \mathbf{K}_{p\eta}\left(\boldsymbol{\eta}_d - \boldsymbol{\eta}_\mathrm{Q}\right) - \mathbf{K}_{d\eta}\boldsymbol{\omega}_\mathrm{Q} \tag{24}$$

and the linearising feedbacks

$$\begin{bmatrix} u_\mathrm{roll} \\ u_\mathrm{pitch} \\ u_\mathrm{yaw} \end{bmatrix} = \begin{bmatrix} \frac{I_x}{K_T L} & 0 & 0 \\ 0 & \frac{I_y}{K_T L} & 0 \\ 0 & 0 & \frac{I_z}{K_Q} \end{bmatrix} \dot{\boldsymbol{\omega}}_d \tag{25}$$

where $\mathbf{K}_{p\eta} = [K_{p\phi}, K_{p\theta}, K_{p\psi}]$ and $\mathbf{K}_{d\eta} = [K_{d\phi}, K_{d\theta}, K_{d\psi}]$.

Table 2: Waypoints for UAV search mode.

Waypoint	x_d	y_d	z_d	ψ_d
1	−1.5	−3.0	−2	$\frac{\pi}{2}$
2	−1.5	3.0	−2	$\frac{\pi}{2}$
3	−1.0	3.0	−2	$-\frac{\pi}{2}$
4	−1.0	−3.0	−2	$-\frac{\pi}{2}$
5	−0.5	−3.0	−2	$\frac{\pi}{2}$
6	−0.5	3.0	−2	$\frac{\pi}{2}$
7	0.0	3.0	−2	$-\frac{\pi}{2}$
8	0.0	−3.0	−2	$-\frac{\pi}{2}$
9	0.5	−3.0	−2	$\frac{\pi}{2}$
10	0.5	3.0	−2	$\frac{\pi}{2}$
11	1.0	3.0	−2	$-\frac{\pi}{2}$
12	1.0	−3.0	−2	$-\frac{\pi}{2}$
13	1.5	−3.0	−2	$\frac{\pi}{2}$
14	1.5	3.0	−2	$\frac{\pi}{2}$

5.3.3. Search Pattern

The search pattern module defines a series of n_w waypoints which the UAV follows after taking off. The waypoint locations are defined such that following each one from the beginning to the end of the search pattern results in near-exhaustive visual coverage of the environment floor. As each waypoint is reached, the counter $c_{\text{wp}} \in \{1, 2, \ldots, n_w\}$ is incremented. Then, if the UAV exits *Search* to investigate a target or return to base, it is able to return to the last waypoint reached. This ensures comprehensive coverage of the environment. The waypoints are listed in Table 2.

5.3.4. Emergency Controller

The emergency controller is similar to the state feedback controller in structure. It is utilised only when an actuator fault occurs and an emergency landing is required. In this event, it is assumed that a single rotor has malfunctioned and it is known which rotor this is. The controller then disables the opposing rotor and increases the thrust to the remaining two rotors to reduce the chance of a hard landing.

With one rotor inoperable, attitude control is neglected to avoid increasing the possibility of a crash. The collective pseudo-input relationship is still given by Equation (22). The remaining pseudo-inputs are fixed at zero. With two rotors inactive, the remaining two rotors must each bear half of the required thrust, giving

$$\mathbf{u}_{\text{Q}} = \begin{cases} [0.5, 0.5, 0, 0]^T u_{\text{col}} & \text{if rotor 3 or 4 is faulty} \\ [0, 0, 0.5, 0.5]^T u_{\text{col}} & \text{if rotor 1 or 2 is faulty} \end{cases} \tag{26}$$

5.3.5. Object Tracking

The object tracking module is used to identify targets within the camera's field of view. The centroid of an identified object in the camera image provides coordinates which are utilised by the visual controller. Thus, upon identifying a target, the UAV is able to accurately position itself above the target and descend to collect it.

The object recognition subsystem emulates that of a simple colour detection algorithm. The targets have distinctive colours which allow them to stand out from each other and the environment. The subsystem masks certain colour ranges in the camera image, thus allowing detection of the targets when they enter the camera field of view. This is emulated in simulation by considering which vertices of the global geometry are within the camera field of view, then identifying the faces and colours associated with these vertices. The actual process of colour detection and the simulated process are described in this section.

Colour Detection Approach. A colour image of resolution $p \times q$ is typically described by a $p \times q \times 3$ array. Different colour models may be used. In this instance, the popular RGB (red-green-blue) model is employed. Each of the three

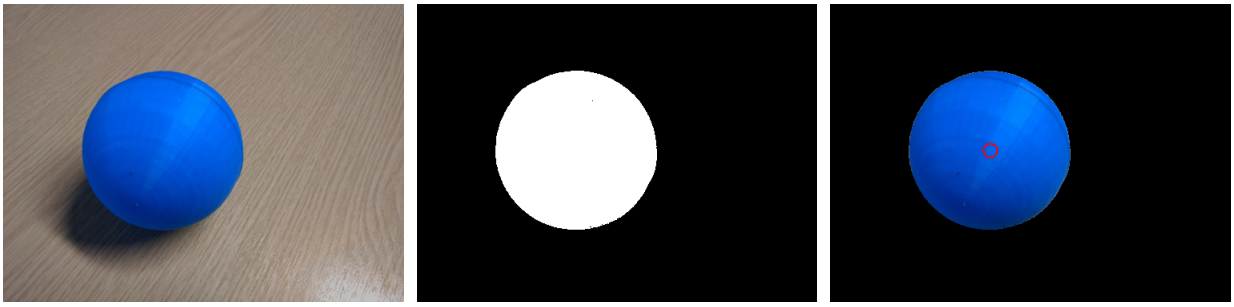

(a) Captured image. (b) Image mask. (c) Mask applied to image, with centroid indicated.

Figure 7: Mask applied to image and centroid indicated.

$p \times q$ slices of the array then describes the intensity of the image in red, green and blue. Intensity values typically vary between 0 and 255. Colour detection is achieved by considering ranges of values for each of the three colours.

Consider the image of a simple monochromatic object, shown in Figure 7a. A mask is generated, which shows only the colours within a certain range, as shown in Figure 7b. The centroid of the unmasked region then provides the approximate geometrical centre of the object, as shown in Figure 7c.

The target objects are defined using distinctive colours which are unique in the environment. The coordinates of the centroid then provide the feedback for the visual controller. This allows the UAV to position itself according to the imagery received by the camera, as interpreted by the colour detection algorithm.

Object Detection Model. The geometry of the model is described in Section 4.3 and the transformation of vertices into camera space described in Section 4.6.3. To be detected, at least one vertex of a single object must satisfy the conditions for visibility to the camera sensor.

Consider the case where m_{vis} vertices are visible to the camera, with coordinates $\mathbf{c}_i, i \in \mathbb{N}^{m_{\text{vis}}} \subseteq \{1, 2, \ldots, m\}$, m being the total number of vertices in the geometrical model of the scenario. That is, each set of coordinates $\mathbf{c}_i$ satisfies the conditions imposed by Equation (14). Each set of coordinates corresponds to a vertex $\mathbf{v}_i$. Each vertex is then associated with one or more faces $\mathbf{f}_j \subseteq \{1, 2, \ldots, m\}$. Thus, each vertex is also associated with one more colours $\mathbf{C}_j \in [0, 1]$, where the RGB range $[0, 255]$ is normalised in the range $[0, 1]$.

Colour detection is achieved by considering only vertices with an associated colour inside a range. For example, to detect red objects, the colour arrays $\mathbf{C}_j$ associated with all visible coordinates $\mathbf{c}_i$ are considered in the limits $\mathbf{C}_{\text{lw}} = [0.4, 0, 0]^T$, $\mathbf{C}_{\text{up}} = [1, 0, 0]^T$. In the simplified geometry we employ in our model, this ignores all colours except the solid red which represents the spherical target. We then consider the b coordinates $\mathbf{c}_i, i \in \mathbb{R}^b$ for which the associated colours satisfy the condition $\mathbf{C}_{\text{lw}} \leq \mathbf{C}_j \leq \mathbf{C}_{\text{up}}$.

Hence, any colour of interest is detected and its vertices identified. We then use a simple centroid calculation to obtain the approximate centre of the target in the camera image

$$\mathbf{c}_{\text{cntrd}} = \frac{1}{b} \sum_{i=1}^{b} \mathbf{c}_i \tag{27}$$

The detection of a target is specified by the boolean $D \in \{0, 1\}$. To avoid detecting targets at the extremes of the image and then losing them due to the UAV motion, a target is only considered as having been identified when its centroid is within a certain radius of the image centre. This boundary is clearly visible in Figure 6. Thus $D = 1$ if $\|\mathbf{c}_{\text{cntrd}}\| < R_c$, where R_c is the radius around the image centre.

Targets already in the drop zone are also ignored. This is accomplished by identifying the drop site position in the camera image $\mathbf{c}_{\text{ds}}$ and rejecting centroids outside the radius $R_{\text{ds,C}}$. Hence $D = 0$ if $\|\mathbf{c}_{\text{ds}}\| < R_{\text{ds,C}}$, where $R_{\text{ds,C}}$ approximates the drop zone radius by

$$R_{\text{ds,C}} = f \frac{R_{\text{ds}}}{|z_{\text{C}}|}$$

The object detection algorithm checks for all three primary colours simultaneously. Where multiple centroids are defined, that is, multiple colours are visible to the camera, the centroid closest to the image centre is supplied to the visual controller. This centroid $\mathbf{c}_{\text{cntrd}}$ is supplied to the visual controller.

5.3.6. Visual Controller

The coordinates of the centroid are denoted $\mathbf{c}_{\text{cntrd}} = [c_x, c_y]$. The position controller acts to drive the coordinates towards the centre of the camera image. To achieve consistency with the state feedback controller, the error is defined as a function of the coordinates and the camera height above the ground, that is

$$\mathbf{e} = \left| \frac{z_{\text{Q}} - z_{\text{C/Q}}}{f} \right| \begin{bmatrix} c_y \\ c_x \end{bmatrix} \tag{28}$$

where $z_{\text{C/Q}}$ is the component of the camera position in $\mathcal{Q}$ and the centroid coordinates are swapped for consistency with the reference frame $\mathcal{Q}$.

A proportional-integral (PI) controller then centres the coordinates in the camera image and the target directly below the UAV. The controller provides a desired velocity command which is simply supplied to the state feedback controller in lieu of the existing horizontal position controller and is given by

$$\begin{bmatrix} \dot{x}_d \\ \dot{y}_d \end{bmatrix} = \begin{bmatrix} \cos\psi & -\sin\psi \\ \sin\psi & \cos\psi \end{bmatrix} \left(K_{pv}\mathbf{e} + K_{iv} \int \mathbf{e}\, dt \right) \tag{29}$$

where K_{pv}, K_{iv} are the proportional and integral gains, respectively.

These commands override the horizontal velocity commands of the state feedback controller in Equation (20). The height of the UAV continues to be controlled by the position controller (Equation (21)) in this mode.

6. Results

The model is used in a Monte Carlo experiment. The stochastic nature of the model has the effect of changing the timing and outcome of events and the ultimate success or failure of the mission. In this section, we first present and discuss the results of a single simulation run. The results of a Monte Carlo simulation conducted over 2,000 runs are then considered. The agent models employ the properties described in Table 3. The UAV's guidance has the parameters described in Table 4.

6.1. Single Run

The simulation is run with landing site, drop site and target positions defined at random and the potential for faults determined by the corresponding probability distribution. In this instance, the mission is successful and takes 170 s to complete.

We can see in Figure 8 the flight path followed by the UAV during the mission. It takes off and flies to the first waypoint at $[-1.5, -3]$ before following a path to the second waypoint at $[-1.5, 3]$. The red spherical target and blue pyramidal target are within both the camera field of view and the object detection algorithm's radius of interest while following this path. The UAV thus diverts to retrieve these targets. While the blue target is identified first, the forward motion of the vehicle results in the red target having closer proximity to the UAV as it enters *Hover above target* mode. The red target is thus retrieved and deposited first.

The UAV then returns to the first waypoint and again follows the path to the second. This ensures that any additional targets along this path are detected. The blue target is again identified. It is retrieved and deposited at the drop site. The UAV again returns to the first waypoint and follows the path to the second waypoint. Finding no further targets, it proceeds to the third, fourth then fifth waypoints. While following the path to the sixth waypoint, it detects the green cuboidal target. It retrieves and deposits this target at the drop site. Having successfully deposited all three targets, the UAV returns to the landing site and the mission ends.

We can obtain further information on the mission by analysing the height history of the UAV, shown in Figure 9. The UAV clearly begins the mission by taking off to a height of one metre. It then immediately lands again, indicating

Table 3: Quadrotor, target and environment properties.

Property	Symbol	Value	Unit
Camera aspect ratio	A	4/3	–
Quadrotor/floor damping coefficient	c_{Q}	75.5	$\mathrm{N\,s\,m^{-1}}$
Target/floor damping coefficient	c_{T}	20	$\mathrm{N\,s\,m^{-1}}$
Camera focal length	f	400	–
Acceleration due to gravity	g	9.81	$\mathrm{m\,s^{-2}}$
Quadrotor moment of inertia about $x^{\mathcal{Q}}$	I_x	0.03	$\mathrm{kg\,m^2}$
Quadrotor moment of inertia about $y^{\mathcal{Q}}$	I_y	0.03	$\mathrm{kg\,m^2}$
Quadrotor moment of inertia about $z^{\mathcal{Q}}$	I_z	0.04	$\mathrm{kg\,m^2}$
Rotor torque gain	K_Q	120	N
Rotor thrust gain	K_T	0.2	N m
Quadrotor/floor spring coefficient	k_{Q}	3775	$\mathrm{N\,m^{-1}}$
Target/floor spring coefficient	k_{T}	1000	$\mathrm{N\,m^{-1}}$
Rotor moment arm	L	0.2	m
Quadrotor mass	m_{Q}	1.51	kg
Target mass	m_{T}	0.4	kg
Probability of actuator fault in period T_a	P_a	0.01	–
Probability of grasper fault in period T_g	P_g	0.05	–
Probability of system fault during *Initialise*	P_s	0.05	–
Drop zone radius	R_{ds}	0.25	m
Quadrotor effective radius	R_{Q}	0.2	m
Target effective radius	R_{T}	0.05	m
Camera position in $\mathcal{Q}$	$\mathbf{r}_{\mathrm{C/Q}}$	$[0, 0, 0.1]^T$	m
Grasper position in $\mathcal{Q}$	$\mathbf{r}_{\mathrm{G/Q}}$	$[0, 0, 0.2]^T$	m
Time period for actuator fault probability P_a	T_a	60	s
Time period for actuator fault probability P_g	T_g	60	s
Maximum battery voltage	V_{max}	11	V
Battery charging rate	v_c	0.025	$\mathrm{V\,s^{-1}}$
Battery discharging rate	v_d	-0.005	$\mathrm{V\,s^{-1}}$
Time constant of gimbal response	τ_g	0.005	s
Camera field of view	λ	45	°

Table 4: Guidance system and controller properties.

Property	Symbol	Value	Unit
Maximum roll/pitch command	$a_{\max}$	20	°
Velocity controller gain	K_{dr}	3.9	–
Roll/pitch rate controller gain	$K_{d\phi}$, $K_{d\theta}$	39	–
Yaw rate controller gain	$K_{d\psi}$	1.95	–
Visual controller integral gain	K_{iv}	0.0001	–
Position controller gain	K_{pr}	0.975	–
Visual controller proportional gain	K_{pv}	0.341	–
Roll/pitch angle controller gain	$K_{p\phi}$, $K_{p\theta}$	380.25	–
Yaw angle controller gain	$K_{p\psi}$	0.951	–
Velocity filter bandwidth	N	50	rad s^{-1}
Maximum velocity command	$v_{\max}$	5	m s^{-1}
Hover height	z_{hvr}	-1	m
Search height	z_{srch}	-2	m
Transport height	z_{trnsprt}	-1	m

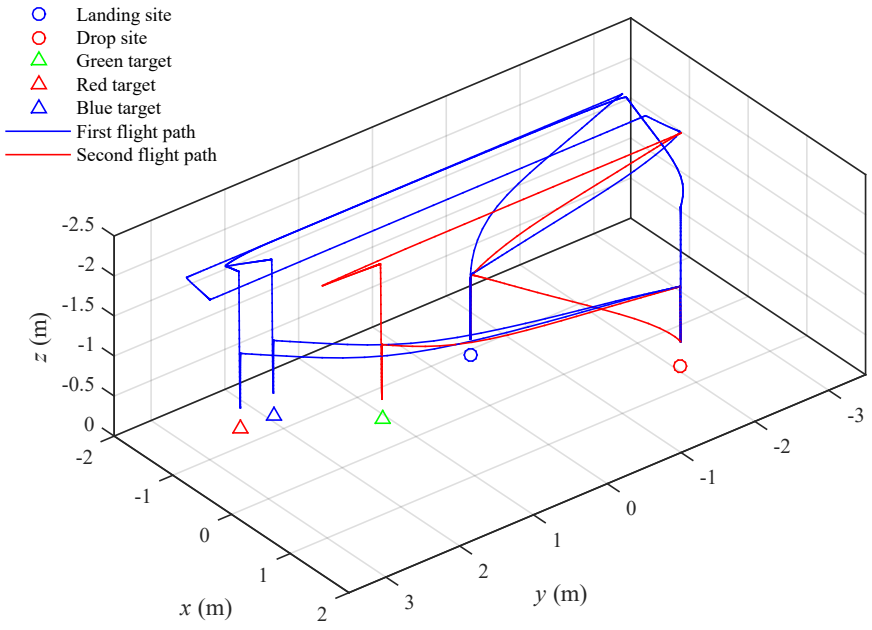

Figure 8: UAV flight path and target position during example mission.

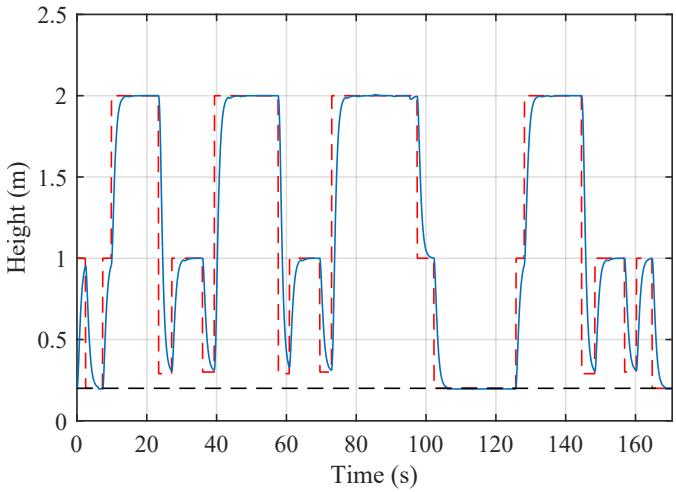

Figure 9: Height history during the mission. The dashed line indicates the height command.

that a system fault has occurred, that is $(1 - P_s) > X \sim \mathcal{U}(0, 1)$. The UAV then takes off again, before ascending to the search height of two metres and continuing the mission.

We can see that the UAV descends multiple times during the mission to an altitude of around 0.3 m. This is the height at which it grasps and drops the targets. It thus occurs six times, grasping and dropping each target. An additional descent occurs at around 100 s. This is followed by a period of around 20 s where the UAV rests on the floor. This is clearly an instance of a low battery warning triggering a return to base, where the UAV sits idle until it has recharged. It then resumes the mission. We may confirm this assumption by checking the battery voltage level history in Figure 10. The voltage level clearly crosses the threshold just prior to the 100 s mark. The UAV then returns to base and recharges the battery until full, before resuming the mission. The mission is then completed before a second recharge is required.

6.2. Monte Carlo Simulation

We may approximate the probability of the mission being successful using Monte Carlo simulation. This provides a result which may be compared to the results of the model checking approach in future work. We may also analyse the paths taken by the finite state machine as the UAV performs the mission.

For the Monte Carlo simulation, we employ a simplified version of the simulation. The quadrotor initial position is fixed at $\mathbf{r}_{Q,0} = [-1.8795, -2.5936, -0.2000]^T$ m with heading $\psi_{Q,0} = 0.9305$ rad. The drop site location is similarly fixed at $\mathbf{r}_{ds} = [1.3303, -1.3307, 0]^T$ m. The number of targets has been reduced to two. The initial positions of the targets are still randomly defined as previously described. The faults which may occur are still subject to the probabilities previously described. These simplifications are employed to facilitate comparison with the results of model checking the scenario in PRISM, which is currently in an earlier state of development.

The simulation is run for 2,000 iterations. The stochastic and probabilistic properties of the model are shown to impact a number of general mission properties. These properties include: the outcome of the mission; the duration of the mission; the times spent and distance travelled in each mode and; the transitions from one mode to another.

6.2.1. Analyses of Mission Failures

The probability of mission success is given in Table 5 as almost 80%. We may examine missions for which the outcome was failure in order to obtain a number of reasons for these failures. First, at least one target may already be in the drop zone as the mission begins. Since the UAV ignores objects in the radius around the drop site, it is unable to find this target and reaches the end of the search path having relocated the remaining targets. Since *all* targets must be retrieved for the mission to be considered successful, this instance results in a mission failure.

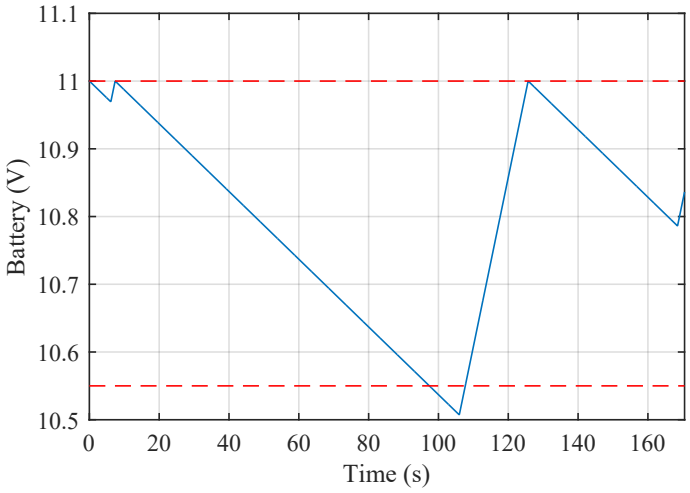

Figure 10: Battery voltage level during the mission. The upper and lower dashed lines indicate the maximum voltage level and warning threshold, respectively.

Table 5: Results of Monte Carlo method.

Parameter	Occurrence
Mission success	0.8750
System fault in *Initialise* mode	0.0484
Actuator fault during mission	0.0335
Grasping fault during target transportation	0.0105

A similar outcome occurs when a target lands in the drop zone but does not register as having been deposited. This occurs when the target is dropped during *Transport* and lands in the drop zone, rendering it invisible to the UAV. The premature release of the target can occur due to a grabber fault or a low battery warning.

Yet another route to mission failure is identified in the random positioning of the targets. One or more targets may be located such that they are ignored by the UAV, despite being outside of the drop zone. This may occur when targets are located in the corners of the environment. The target is outside of the object tracking algorithm's radius of interest R_c and, without any overlapping fly-bys occurring in this area, the target is ignored.

The remainder of mission failures occur due to actuator faults, which may occur at any time. Since returning to the landing site is a prerequisite for mission success, an emergency landing results in mission failure. This occurs regardless of whether all three targets have been retrieved and deposited at the drop site.

6.2.2. Probabilities Derived from Simulation

The probabilities of each of the three faults occurring are also given in Table 5. A system fault may occur during *Initialise* only. The probability of this occurring was given the explicit value $P_s = 0.05$. We use this probability as a verification check, with the result that the measured occurrence of a system fault in *Initialise* is $P_{s,m} = 0.0484$. The result is found from over 4,000 samples, thus we can state with some confidence that a system fault occurs with expected frequency.

The actuator fault is specified as having a certain probability $P_a = 0.01$ of occurring with a given time $T_a = 60$ s. The results of the Monte Carlo experiment provide the occurrence of an actuator fault as percentage of the total number of iterations. As the actuator fault results in the mission ending in failure, it may only occur once per mission. An actuator fault is found to occur in 3.35% of missions.

The grasping mechanism fault is similarly defined as having a probability $P_g = 0.05$ of occurring in a $T_g = 60$ s period. This fault can occur at any time during the mission. However, it only has impact when the mechanism is active, that is, when a target is tethered to the UAV. The occurrence of a grasping fault is thus limited to those modes during which a target is tethered to the UAV. A grasping mechanism fault is then found to cause a premature target drop in 1.05% of the missions.

Mission time is found to have an average of 137.2 s.

We may also consider the probabilities of progressing from one mode to another. The finite state machine in Figure 5 may be redrawn as Figure 11, where the predicates for the transitions are replaced with the probabilities of a given transition occurring. Multiple modes may lead to both *Emergency land* and *Return to base*. The probabilities of these transitions occurring are listed separately in Table 6 for clarity.

In the majority of cases, it is clear that there is a favoured mode to progress to from a given initial mode. Consider the transition from *Initialise* to *Search*, which has a much higher probability of occurring that a transition from *Initialise* to *Land*. In contrast, while *Hover above target* is most likely to progress to *Descend to grasp*, the probability of the UAV instead proceeding to *Search* is non-negligible.

As expected, the probability of progressing from *Drop* to *Return to base* versus *Return to search* is near equal. Since there are only two targets, the UAV will either *Return to search* after dropping the first target or *Return to base* after dropping the final target.

7. Symbolic Abstract Model

As discussed in the introduction, the simulation model developed in this paper is intended to aid in the development of a framework for analysing autonomous systems and, in particular, the synthesis of high-level control for such systems. As an initial investigation, we are using the simulation model in the development and validation of a symbolic abstract model of the same scenario [31]. In the symbolic model the quantitative data for abstract modes and actions is derived from simulating small-scale models which, in turn, are derived from the presented simulation model. When building the abstraction non-determinism is introduced as the precise behaviour of the system is lost. The symbolic model is therefore represented as a Markov Decision Process (MDP) which exhibits both probabilistic and non-deterministic choice [32].

An MDP consists of a set of states, set of actions, probabilistic transition function and reward function. In any state of an MDP there is a set of available actions (a subset of the action set) and for any state and available action the

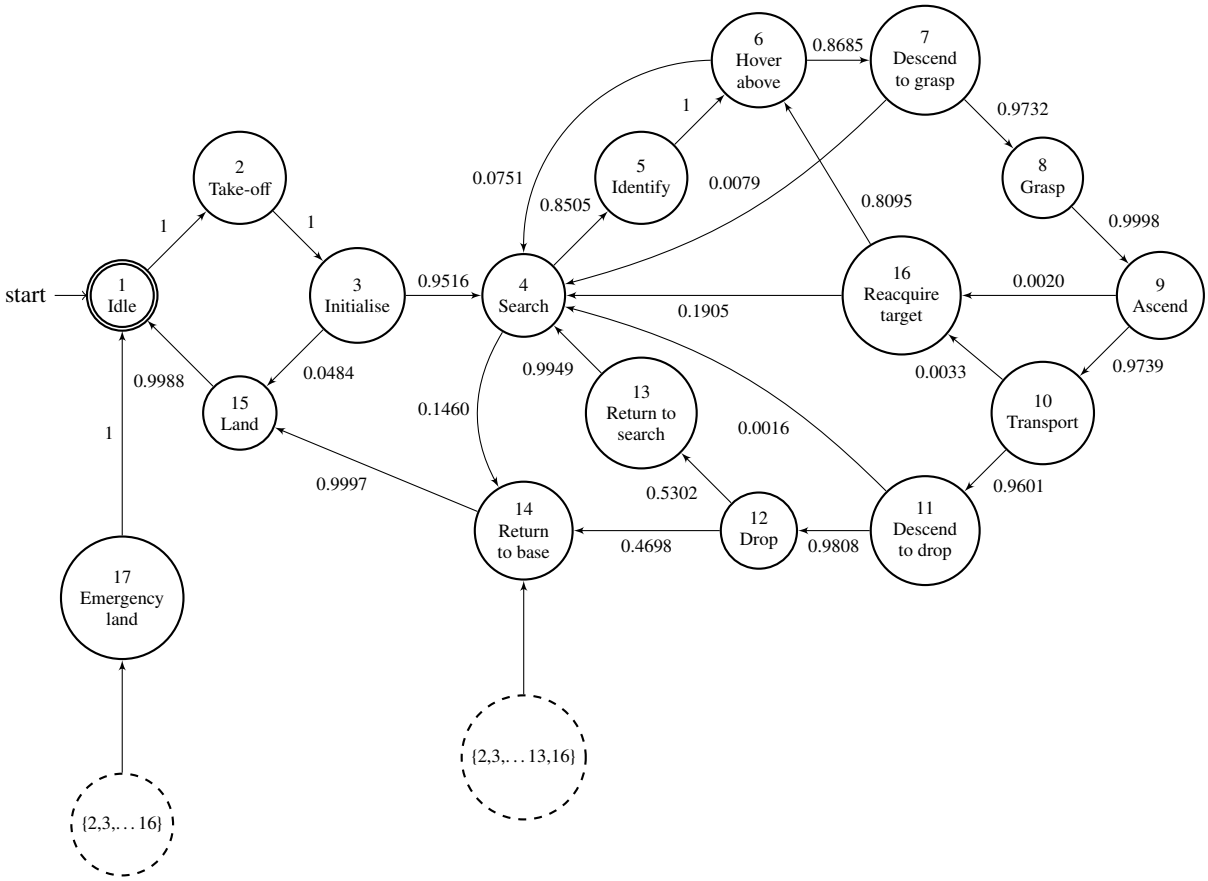

Figure 11: Finite-state machine showing probabilities of transitioning from one state to another.

Table 6: Probabilities of transitioning from a given state/mode to *Emergency land* or *Return to base*.

Mode	Emergency land	Return to base
Idle	0	0
Take-off	0	0
Initialise	0	0
Search	0.0036	0.1460
Identify	0	0
Hover above target	0.0025	0.0539
Descend to grasp	0.0034	0.0156
Grasp	0	0.0002
Ascend	0.0017	0.0224
Transport	0.0010	0.0356
Descend to drop	0	0.0176
Drop	0	0.4698
Return to search	0	0.0051
Return to base	0.0003	–
Land	0.0012	–
Reacquire target	0	–
Emergency land	–	–

probabilistic transition function returns a distribution over successor states and the reward function returns the reward accumulated when taking the action. The behaviour in any state consists of a non-deterministic choice from the set of available actions, followed by a probabilistic choice made according to the probabilistic transition function. An execution of an MDP is an sequence of such choices. The non-deterministic choices of available actions in each state are resolved by a strategy (also called an adversary or scheduler) which can base its choices on the execution up to that point. For a given strategy, the behaviour of an MDP is captured by a probability measure over the executions corresponding to the strategy's choices.

Key properties for MDPs are the probability of reaching a target set of states and the expected reward cumulated until this occurs. Of particular interest are the minimum and maximum values for these properties over all strategies. In models where the non-determinism in the MDP is through the abstraction of a concrete value, the minimum and maximum values yield upper and lower bounds on the precise value for the concrete system [33].

7.1. Model Checking with PRISM

We have built and analysed the symbolic model using the probabilistic model checker PRISM [3, 34]. Models in PRISM are expressed using a high level modelling language based on the Reactive Modules formalism [35]. A model consists of a number of interacting modules. Each module consists of a number of finite-valued variables corresponding to the module's state and the transitions of a module are defined by a number of guarded commands of the form:

$$[\texttt{<action>}]\ \texttt{<guard>} \rightarrow \texttt{<prob>} : \texttt{<update>} + \cdots + \texttt{<prob>} : \texttt{<update>}$$

A command consists of an (optional) action label, a guard and probabilistic choice between updates. A guard is a predicate over the variables of the modules, while an update specifies, using primed variables, how the variables of the module are updated when the command is taken. Interaction between modules is through both the guards (as guards can refer to variables of other modules) and action labels which allow modules to synchronise. Support for reward functions are through reward items of the form:

$$[\texttt{<action>}]\ \texttt{<guard>} : \texttt{<reward>};$$

representing the reward accumulated when taken taking an action in a state satisfying guard. A reward is a real-valued expression which can include variables and constants of the model.

Our PRISM specification consists of four modules: a *UAV* module, *Time/Battery/Actuator* module, *Movement* module and *Object* module. The complete PRISM code can be found in the repository [36]. To illustrate the abstract PRISM model, we present some fragments of the PRISM code corresponding to abstract actions representing search, grasp and release. In these code fragments the following variables appear:

- *s1* representing the current mode of the UAV;
- *posx* and *posy* representing the coordinates of UAV (the height of the UAV is abstracted);
- *retx* and *rety* representing the position of the last way point;
- *c* representing the status of the actuator (operational or failed);
- *b* representing the battery charge level;
- *t* representing the current mission time;
- *NoOfObjs* representing the number of objects still to be found.

Fragments of the PRISM code for the abstract action search are presented in Figure 12. In the *UAV* module, the UAV stays in the search mode if it has not found an object and has not reached the last possible coordinates (of the arena). If an object has been found, there is a transition to the identify mode.

In the *Time/Battery/Actuator* module, in accordance with the simulation model and small scale simulation results, the UAV moves one grid square per second, the rate of depletion of the battery is one unit per second and the probability of the actuator failing in one second is 0.00018. If the actuator does not fail (c=0), the battery level and time are

```
// formula to check if search has found either object 1 or 2
formula FOUND = ((posx=Objx1 & posy=Objy1) | (posx=Objx2 & posy=Objy2));

// UAV Module - search mode (s1=3)
   // found object (record way point for returning to search) and move to identify mode
   [srch] s1=3 & FOUND → (s1'=4)&(retx'=posx)&(rety'=posy);
   // not found object and not finished search (continue in search mode)
   [srch] s1=3 & !(posx=lastX & posy=lastY) & !FOUND → true;
   // not found object and finished search (move to return to base mode)
   [srch] s1=3 & posx=lastX & posy=lastY & !FOUND → (s1'=12)&(retx'=posx)&(rety'=posy);

// Time/Battery/Actuator Module - search
// Actuator fault occurs in one time step with probability pf
// if operational update time and charge, otherwise update actuator
   // continue search - action can be done within time limit and battery is sufficiently charged
   [srch] c=0 & b>Blow & t+dt<Miss → pow(1−pf, dt) : (t'=t+dt)&(b'=b−db∗dt) + (1−pow(1−pf, dt)) : (c'=1);

// Movement Module - search
   // moving along the x-axis to the right
   [srch] posx≥0 & posx<Xcoord & posy≤Ycoord & mod(posy, 2)=0 & !FOUND → (posx'=posx+1);
   // moving along the y-axis (reached the right boundary of the arena)
   [srch] posx=Xcoord & mod(posy, 2)=0 & posy<Ycoord & !FOUND → (posy'=posy+1);
   // moving along the x-axis to the left
   [srch] posx>0 & posx≤Xcoord & posy≤Ycoord & mod(posy, 2)=1 & !FOUND → (posx'=posx−1);
   // moving along the y-axis (reached left boundary of the arena)
   [srch] posx=0 & mod(posy, 2)=1 & posy<Ycoord & !FOUND → (posy'=posy+1);
   // found an object
   [srch] FOUND → true;
   // reached the last search coordinates
   [srch] posx=lastX & posy=lastY → true;
```

Figure 12: PRISM code relating to the abstract action search.

```
// Time constants
const int Tgl = 0; // time to grab/release lower bound
const int Tgu = 1; // time to grab/release upper bound

// UAV Module - grasp mode (s1=6), drop mode (s1=10)
   // grab object and start ascending moving to base
   [grab] (s1=6) → (s1'=7);
   // drop object and start ascending (towards search or base)
   [rel] (s1=10) → (s1'=11);

// Time/Battery/Actuator Module
// Actuator fault occurs in one time step with pf probability
// if operational update time and charge, otherwise update actuator
   // grab - upper time bound - action can be done within time limit and battery is sufficiently charged
   [grab] c=0 & b>Blow & t+Tgu<Miss → pow(1−pf, Tgu) : (t'=t+Tgu)&(b'=b−db∗Tgu) + (1−pow(1−pf, Tgu)) : (c'=1);
   // grab - lower time bound - action can be done within time limit and battery is sufficiently charged
   [grab] c=0 & b>Blow & (t+Tgl<Miss) → pow(1−pf, Tgl) : (t'=t+Tgl)&(b'=b−db∗Tgl) + (1−pow(1−pf, Tgl)) : (c'=1);
   // release - upper time bound - action can be done within the time limit and battery is sufficiently charged
   [rel] c=0 & b>Blow & t+Tgu<Miss → pow(1−pf, Tgu) : (t'=t+Tgu)&(b'=b−db∗Tgu) + (1−pow(1−pf, Tgu)) : (c'=1);
   // release - lower time bound - action can be done within the time limit and battery is sufficiently charged
   [rel] c=0 & b>Blow & t+Tgl<Miss → pow(1−pf, Tgl) : (t'=t+Tgl)&(b'=b−db∗Tgl) + (1−pow(1−pf, Tgl)) : (c'=1);

// Object Module
   // grab object from where it was found
   [grab] s1!=10 → true;
   // release - object has been deposited safely, one less object to search for
   [rel] s1=10 & NoOfObjs>0 → (NoOfObjs'=NoOfObjs−1);
```

Figure 13: PRISM code fragments relating to the grasp and release abstract actions.

updated. In the abstract model when a transition corresponds to more than one second elapsing we have to ensure we have the correct values for the battery depletion and the cumulative probability of failure. More precisely, if a transition in the abstract model takes t seconds, then when taking this transition the battery gets depleted by t units and the probability that the actuator does not fail is $(1-pf)^t$ (and hence the probability of failure is $1-(1-pf)^t$).

```
// lower and upper bounds for the probability of a successfully mission
Pmin=? [ F "MissionSuccessful" ];
Pmax=? [ F "MissionSuccessful" ];
// lower and upper bounds for an actuator fault occuring
Pmin=? [ F "fault" ];
Pmax=? [ F "fault" ];
// lower and upper bounds for the expected mission time
R{"time"}min=? [ F "done" ];
R{"time"}max=? [ F "done" ];
```

Figure 14: PRISM properties code.

The fragments of the PRISM code for the three modules relevant to grab and release actions are presented in Figure 13. The grasp and release abstract actions were both shown to take 0.01 seconds via small-scale simulation. Through non-determinism (see Figure 13), the PRISM model abstracts these timing values to lower and upper bounds (0 and 1). In the *Object* module the number of non-deposited objects *NoOfObjs* is decremented when an object is released at the depot site.

We evaluated the symbolic implementation by building a model for each combination of object positioning in the arena, while keeping the base and the depot sites fixed. The properties that we focused on determine probability bounds for a mission to be successful, and for a mission to end due to an actuator fault respectively. The PRISM properties are shown in Figure 14 and we have computed minimum and maximum values yielding upper and lower bounds for the concrete system. The bounds obtained are:

- $[0.661, 0.891]$ for the probability of successful mission completion (the simulation results returned 0.8750);
- $[0.018, 0.041]$ for the probability of an actuator fault (the simulation results returned 0.0335);
- $[101.3, 223.7]$ for the expected mission time (the simulation results returned 137.2).

As can be seen the bounds from the symbolic model encompass the outcomes from the Monte Carlo simulations presented in Section 6. We can also see that the relative differences of the bounds are of the same order of magnitude. Although the bounds are not tight, they are promising for an initial investigation. One potential reason for the bounds not being tight is how time has been abstracted in the model. In particular, in some cases the lower bound on the time taken for some transitions is 0 (as can be seen in Figure 13), allowing certain abstract actions to be taken repeatedly without influencing the overall mission time. This behaviour is clearly not realisable in the simulation model.

8. Conclusions

We may draw some conclusions on the verification of real-world autonomous systems from both the development of our continuous-time model and the results of the simulation.

8.1. Limitations of the Model

As an entity in simulation only, our model simplifies or omits some complex real-world behaviours. First, target impact on UAV behaviour is not modelled. In reality, grasping the target would alter the mass and moments of inertia of the UAV. This would then affect the closed-loop response of the vehicle, potentially changing the outcome of the mission. Specifically, the impact of the target mass on the state feedback controller must be considered in any practical implementation of the scenario.

Additionally, no fault detection algorithm is implemented. Faults in the system are assumed to be detected instantaneously. The autonomous controller then switches to an appropriate mode to mitigate the effects of the fault. In reality, the faults would be detected either be sensors monitoring the affected components or by any simple fault detection method [37].

The lack of disturbances and the use of feedback linearisation in the controller precludes the need for integral action. This is because the linearising feedbacks perfectly cancel out the non-linearities and system properties in

simulation [29]. In reality, this feedbacks do not perfectly cancel out the system dynamics. Thus, there is a need for a more robust controller. The presence of disturbances in reality further establishes the need for integral action.

While a simple power model is implemented in our quadrotor agent, its role is simply as an additional predicate for a greater variety of transitions in the finite state machine. In reality, a decreasing battery voltage may impact the performance of the rotors and can affect the flight characteristics of the UAV even before a return to base command is issued.

The camera sensor and grasping mechanism are optimally positioned such that an object centred in the camera image is directly below the grasping mechanism. In reality, the grasper would obscure the object from view. In implementing the autonomous quadrotor system, we would need to carefully consider the positions of the camera and grasper and their impact on the target detection and grasping.

Finally, we have simplified the interaction and operation of the UAV's subsystems by assuming they function flawlessly. For example, the camera sensor is assumed to have a high framerate and perceives the targets as solid blocks of colour. The object detection module is assumed to execute at a high bandwidth, resulting in no lag between the camera and the eventual centroid coordinates. In reality, the object detection will likely be imperfect and will run at a relatively low bandwidth to ensure real-time execution. This would then impact both the visual controller and the autonomous mode switching.

8.2. Limitations of the Autonomous Controller

The modes of the autonomous controller FSM were primarily determined prior to implementation. Early simulation testing resulted in the occurrence of events for which the FSM had not been programmed to consider. Some additional modes were thus incorporated into the FSM, such as *Reacquire target* or *Return to search*. Experimental testing in the future is likely to highlight further gaps in the FSM and require the addition of more guidance modes.

The current implementation of the autonomous controller is deliberately lacking in parts, to better facilitate the chance of mission failures. Analysis of the mission specifics for each of the 2,000 iterations provides a number of reasons for mission failure.

First, there is the possibility that the UAV believes the mission has failed when it has in fact been successful. This may occur when the location of one of the targets is randomly defined to be already in the specified radius around the drop site. The UAV is unable to detect this target. It therefore follows the mission in depositing the remaining targets, before reaching the end of the search path and returning to base, believing the mission has failed. This outcome is a consequence of the fact that both target location and drop site are defined randomly and may therefore coincide. The autonomous behaviour of the UAV does not conceive of this occurrence, knowing only that there are three targets and that it has deposited two.

A similar outcome can occur when all targets are initially outside of the drop site radius, but at least one is somehow undetectable to the UAV. This only occurs when the target is positioned in the corner of the environment, as the overlap in sensor perception elsewhere prevents this occurrence otherwise.

8.3. Future Work

We aim to continue this work on two fronts. First, we are continuing our work in formal verification of autonomous systems in PRISM. The symbolic model of our system must satisfy the dual requirements of accurately describing the intricacies of the system, while also being sufficiently simple that it may be verified in a short period of time. The pursuit of this research is intended to lead to the creation of a framework for the verification of autonomous systems. For that we need to formally prove a connection between the simulation and symbolic models which will allow us to infer results for the actual system from those obtained through model checking the symbolic model. The longer term goal is to develop a highly adaptable framework for analysing autonomous systems. For example, once the initial small-scale simulation models for abstract actions and symbolic models have been set up for a certain scenario, different decision algorithms can be easily substituted and analysed. In addition, we can reuse the quantitative data for abstract actions obtained from the small scale simulation models when analysing different scenarios in which the same abstract actions occur.

On the second front, we are implementing our autonomous mission in controlled laboratory conditions. By doing this, we hope to identify additional factors which affect the mission outcome and determine whether these factors may be mitigated by robust system design or by real-time system verification. We are also investigating the required hardware to support real-time verification and onboard autonomy.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/N508792/1].

References

[1] M. Webster, M. Fisher, N. Cameron, and M. Jump. Formal Methods for the Certification of Autonomous Unmanned Aircraft Systems. *Computer Safety, Reliability, and Security*, 6894:228–242, 2011.

[2] F. Lerda, J. Kapinski, H. Maka, E. Clarke, and B. Krogh. Model Checking In-The-Loop. In *Proceedings of the American Control Conference*, pages 2734–2740, 2008.

[3] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In G Gopalakrishna and S Qadeer, editors, *Proceedings of the 23rd International Conference on Computer Aided Verification (CAV'11)*, volume 6806 of *LNCS*, pages 585–591. Springer, Jul 2011.

[4] H Abbas, G Fainekos, S Sankaranarayanan, F Ivancic, and A Gupta. Probabilistic Temporal Logic Falsification of Cyber-Physical Systems. *ACM Transactions on Embedded Computing Systems*, 12(2):30, 2013.

[5] S. Sankaranarayanan and G. Fainekos. Falsification of Temporal Properties of Hybrid Systems Using the Cross-Entropy Method. In *Proceedings of the 15th International Conference on Hybrid Systems: Computation and Control*, pages 125–134, Beijing, apr 2012.

[6] T. Nghiem, S. Sankaranarayanan, G. Fainekos, F. Ivancic, A. Gupta, and G. Pappas. Monte-carlo techniques for falsification of temporal properties of non-linear hybrid systems. In *Proceedings of the 13th International Conference on Hybrid Systems: Computation and Control*, pages 211–220, Stockholm, apr 2010.

[7] M. Althoff and . Dolan. Online verification of automated road vehicles using reachability analysis. *IEEE Transactions on Robotics*, 30(4):903–918, 2014.

[8] J. Gillula, H. Huang, M. Vitus, and C. Tomlin. Design of guaranteed safe maneuvers using reachable sets: Autonomous quadrotor aerobatics in theory and practice. In *Proceedings - IEEE International Conference on Robotics and Automation*, pages 1649–1654, 2010.

[9] D. Althoff, M. Althoff, and S. Scherer. Online Safety Verification of Trajectories for Unmanned Flight with Offline Computed Robust Invariant Sets. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3470–3477, Hamburg, sep 2015. IEEE.

[10] T. Bochot, P. Virelizier, H. Waeselynck, and V. Wiels. Model Checking Flight Control Systems: The Airbus Experience. In *31st International Conference on Software Engineering - ICSE '09*, pages 18–27, Vancouver, May 2009. IEEE.

[11] T. Henzinger. The Theory of Hybrid Automata. In *Proceedings 11th Annual IEEE Symposium on Logic in Computer Science*, pages 278–292, New Brunswick, NJ, Jul 1996. IEEE.

[12] A. Chutinan and B. Krogh. Verification of Infinite-State Dynamic Systems Using Approximate Quotient Transition Systems. *IEEE Transactions on Automatic Control*, 46(9):101–109, Sep 2000.

[13] S Kowalewski, S Engell, J Preußig, and O Stursberg. Verification of logic controllers for continuous plants using timed condition/event-system models. *Automatica*, 35(3):505–518, 1999.

[14] C. Bohn. Heuristics for Designing the Control of a UAV Fleet with Model Checking. *Lecture Notes in Economics and Mathematical Systems*, 588:21–36, 2007.

[15] S. Konur, C. Dixon, and M. Fisher. Analysing robot swarm behaviour via probabilistic model checking. *Robotics and Autonomous Systems*, 60(2):199–213, 2012.

[16] W. Liu, A.T Winfield, and J. Sa. Modelling Swarm Robotic Systems : A Case Study in Collective Foraging. In *Proceedings of Towards Autonomous Robotic Systems*, pages 25–32, Aberystwyth, Sep 2007.

[17] L. Mikaël. *Formal Verification of Flexibility in Swarm Robotics*. Masters thesis, Université libre de Bruxelles, 2012.

[18] L. Humphrey. Model Checking for Verification in UAV Cooperative Control Applications. In Fariba Fahroo, Le Yi Wang, and George Yin, editors, *Recent Advances in Research on Unmanned Aerial Vehicles*, chapter 4. Springer, 2013.

[19] J. Choi. *Model Checking for Decision Making Behaviour of Heterogeneous Multi-Agent Autonomous System*. Phd thesis, Cranfield University, 2012.

[20] Y. Kubera, P. Mathieu, and S. Picault. Everything can be Agent! (Extended Abstract). In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*, pages 1547–1548, Toronto, May 2010.

[21] A. Frank, S. Bittner, and M. Raubal. Spatial and Cognitive Simulation with Multi-agent Systems. In Daniel R Montello, editor, *Spatial Information Theory: Foundations of Geographic Information Science*, pages 124–139. Springer Berlin Heidelberg, 2001.

[22] S. Bouabdallah, P. Murrieri, and R. Siegwart. Design and Control of an Indoor Micro Quadrotor. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 4393–4398. IEEE, 2004.

[23] H. Voos. Nonlinear Control of a Quadrotor Micro-UAV Using Feedback-Linearization. In *Proceedings of the 2009 IEEE International Conference on Mechatronics*, Malaga, 2009. IEEE.

[24] M. Ireland. *Investigations in Multi-Resolution Modelling of the Quadrotor Micro Air Vehicle*. PhD, University of Glasgow, may 2014.

[25] A. Tayebi and S. McGilvray. Attitude Stabilization of a VTOL Quadrotor Aircraft. *IEEE Transactions on Control Systems Technology*, 14(3):562–571, may 2006.

[26] M. Mueller and R. D'Andrea. Stability and control of a quadrocopter despite the complete loss of one, two, or three propellers. In *Proceedings of the 2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 45–52, Hong Kong, may 2014. Ieee.

[27] R Hartley and A Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, 2nd edition, 2003.

[28] M. Ireland and D. Anderson. An Investigation of the Effects of Model Resolution on Control of a Quadrotor Micro Air Vehicle. *IJUSEng*, 3(1):17–25, 2015.

[29] M. Ireland, A. Vargas, and D. Anderson. A Comparison of Closed-Loop Performance of Multirotor Configurations Using Non-Linear Dynamic Inversion Control. *Aerospace*, 2(2):325–352, 2015.

[30] A. Das, K. Subbarao, and F. Lewis. Dynamic inversion with zero-dynamics stabilisation for quadrotor control. *Control Theory & Applications, IET*, 3(3):303–314, 2009.

[31] R. Hoffmann, M. Ireland, A. Miller, G. Norman, and S. Veres. Autonomous Agent Behaviour Modelled in PRISM: A Case Study. In Bonaki, D and Wijs, A, editor, *23rd International SPIN Symposium on Model Checking of Software*, volume 9641 of *LNCS*, Eindhoven, April 2016. [Accepted for publication].

[32] M. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, 1994.

[33] V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker. Automated verification techniques for probabilistic systems. In M. Bernardo and V. Issarny, editors, *Formal Methods for Eternal Networked Software Systems (SFM'11)*, volume 6659 of *LNCS*, pages 53–113. Springer, 2011.

[34] PRISM web site. www.prismmodelchecker.org.

[35] R. Alur and T. Henzinger. Reactive modules. *Formal Methods in System Design*, 15(1):7–48, 1999.

[36] Abstract PRISM model. BitBucket Repository, 2016. bitbucket.org/ruthhoffmann/symbolic-uav-model.

[37] R. Isermann. *Fault-Diagnosis Systems*. Springer, 2006.